OAK RIDGE NATIONAL LABORATORY

MANAGED BY UT-BATTELLE FOR THE DEPARTMENT OF ENERGY

US Department of Energy, Office of Science

High Performance Computing Facility Operational Assessment 2013 Oak Ridge Leadership Computing Facility

February 2014

Prepared by

Ashley D. Barker
David E. Bernholdt
Arthur S. Bland
James J. Hack
Douglas L. Hudson
James H. Rogers
T.P. Straatsma
Kevin G. Thach
Sudharshan S. Vazhkudai
Jack C. Wells
Julia C. White

UT-BATTELLE

ORNL-27 (4-00)

US Department of Energy, Office of Science

HIGH PERFORMANCE COMPUTING FACILITY OPERATIONAL ASSESSMENT 2013
OAK RIDGE LEADERSHIP COMPUTING FACILITY

Ashley D. Barker
David E. Bernholdt
Arthur S. Bland
James J. Hack
Douglas L. Hudson
James H. Rogers
T.P. Straatsma
Kevin G. Thach
Sudharshan S. Vazhkudai
Jack C. Wells
Julia C. White

February 2014

Prepared by
OAK RIDGE NATIONAL LABORATORY
Oak Ridge, Tennessee 37831-6283
Managed by
UT-BATTELLE, LLC
for the
US DEPARTMENT OF ENERGY
under contract DE-AC05-00OR22725

CONTENTS

Page

LIST OF FIGURES

Figure **Page**

LIST OF TABLES

ABBREVIATED TERMS

ACCEL	Accelerating Competitiveness through Computational Excellence
ACME	Accelerated Climate Modeling for Energy
ACS	American Chemical Society
ACSS	Accelerating Computational Science Symposium
ACTS	Assessment and Commitment Tracking System
ALCC	ASCR Leadership Computing Challenge
ALCF	Argonne Leadership Computing Facility
ALTD	Automatic Library Tracking Database
ANL	Argonne National Laboratory
ASCR	Advanced Scientific Computing Research
ATO	authority to operate
BNL	Brookhaven National Laboratory
BoF	Birds of a Feather
C&A	certification and accreditation
CAAR	Center for Accelerated Application Readiness
CASL	Center for Advanced Simulation of Light Water Reactors
CDL	California digital Library
CFD	computational fluid dynamics
CFR	Code of Federal Regulations
CNMS	Center for Nanophase Materials Sciences
CCSD	Computing and Computational Sciences Directorate
DAC	Directorate Advisory Committee
DD	Director's Discretionary
DNE	Distributed Namespace
DOE	US Department of Energy
DOI	digital object identifier
EERE	Office of Energy Efficiency and Renewable Energy
EGR	exhaust gas recirculation
ES	Early Science
EVEREST	Exploratory Visualization Environment for Research in Science and Technology
FDR	fourteen data rate
FIT	failure in time
FPU	floating point unit
FY	fiscal year
GPU	graphics processing unit
HACC	Hardware/Hybrid Accelerated Cosmology Code
HIS	high-speed interconnect
HPC	high-performance computing
HPL	High-Performance Linpack
HPSS	High-Performance Storage System
HQ	headquarters
IG	inspector general
INCITE	Innovative and Novel Computational Impact on Theory and Experiment
ISI	Institute for Scientific Information
ISM	Integrated Safety Management
IT	information technology
LAMMPS	Large-Scale Atomic/Molecular Massively Parallel Simulator
LANL	Los Alamos National Laboratory

LBNL	Lawrence Berkeley National Laboratory
LDRD	Laboratory Directed Research and Development
LLNL	Lawrence Livermore National Laboratory
LOTO	lock-out/tag-out
LSMS	locally self-consistent multiple scattering code
MPI	Message-Passing Interface
MTTF	mean time to failure
MTTI	mean time to interrupt
NAM	not a metric
NCCS	National Center for Computational Sciences
NERSC	National Energy Research Scientific Computing Center
NIP	not in production
NOAA	National Oceanic and Atmospheric Administration
OA	overall availability
OAR	operational assessment report
OLCF	Oak Ridge Leadership Computing Facility
OMB	Office of Management and Budget
OMP	open multiprocessing
OPV	organic photovoltaic
ORAU	Oak Ridge Associated Universities
ORISE	Oak Ridge Institute for Science and Education
ORNL	Oak Ridge National Laboratory
OSO	DOE ORNL Site Office
OSS	object storage server
OSTI	Office of Scientific and Technical Information
P3HT	poly(3-hexylthiophene)
PCBM	[6,6]-phenyl-C61-butyric acid methyl ester
PFS	parallel file system
PI	principal investigator
PM	preventative maintenance
PTS	Publication Tracking System
R&D	research and development
RFP	request for proposals
RMP	risk management plan
RSS	research safety summary
RT	request tracker
SA	scheduled availability
SASEF	Southern Appalachian Science and Engineering Fair
SBMS	Standards-Based Management System
SC	supercomputing
SciComp	Scientific Computing Group
SIDI	spark-ignited direct-injection
SM	streaming multiprocessor
SNMP	Simple Network Management Protocol
tar	tape archive
UA	User Assistance
ULFM	User-Level Fault Mitigation
UT	The University of Tennessee

Executive Summary

HIGH PERFORMANCE COMPUTING FACILITY
2012 OPERATIONAL ASSESSMENT
OAK RIDGE LEADERSHIP COMPUTING FACILITY

February 2014

EXECUTIVE SUMMARY

Oak Ridge National Laboratory’s (ORNL’s) Leadership Computing Facility (OLCF) leads the way in providing the most powerful computing resource in the United States for open science. In calendar year (CY) 2013 the OLCF accepted the upgrade of the Cray XT Jaguar to the hybrid-architecture Cray XK7 Titan, which delivers a peak performance of more than 27 petaflops. The effectiveness with which this resource was delivered is demonstrated by the business result metrics, which were met or exceeded in all cases. At year’s end, the OLCF delivered all of the compute hours committed to the three major allocation programs: 2.4 billion Titan core-hours were used by Innovative and Novel Computational Impact on Theory and Experiment (INCITE), ASCR Leadership Computing Challenge (ALCC), and Director’s Discretionary (DD) projects. OLCF computational resources support scientific research through production simulation across many scientific domains, providing the key computing and data resources that are critical to their success. Despite the significant disruptions necessary for the Titan upgrade, users continue to express great satisfaction with the OLCF overall.

As leadership systems grow in size, so potentially does the amount of critical data generated during and after the simulations. These data are, along with publications, the intellectual capital of the research communities. The OLCF continues to focus attention on ways to facilitate scientific accomplishments through efforts to improve generation and movement of data, its access and analysis. In 2013 the OLCF formalized the addition of data liaisons to its science support team and initiated a beta-test of a new DD project type—the data project—as the facility continues to seek to understand the ways in which data creation, storage, curation, exploration, and technical support lead to scientific insight.

OLCF users’ scientific and technical accomplishments are wide-ranging. Several representative highlights are presented in this report and serve to communicate how the OLCF is advancing two of DOE’s four strategic goals, and associated targeted outcomes, as stated in the *U.S. Department of Energy Strategic Plan (May 2011).* User research campaigns resulted in two hundred sixty-two publications in 2013, including articles in *Nature*, *Nature Communications*, and *Nature Scientific Reports*.

Effective operations of the OLCF play a key role in the scientific missions and accomplishments of its users. This Operational Assessment Report (OAR) delineates the policies, procedures, and innovations implemented by the OLCF to continue delivering a multipetaflop resource for cutting-edge research. This report covers CY2013, which, unless otherwise specified, denotes January 1, 2013, through December 31, 2013.

Highlights of OLCF operational activities for 2013 include the following:

- Overall ratings for the OLCF were positive; 95% of users reported being “satisfied” or “very satisfied.” No users reported being “dissatisfied” or “very dissatisfied” with the OLCF overall.

- 59.38% of the delivered compute hours on Titan were the result of capability jobs, significantly above the 2013 target of 30%. This indicates the success of the user communities in using the machine at scale.

- In 2013, the OLCF successfully designed, acquired, and deployed the Spider II parallel file system (PFS). Spider II is four times faster and three times larger than the previous PFS.

- The OLCF upgraded the center's primary visualization resource, EVEREST, and acquired and deployed Rhea, a data analysis and visualization cluster and a four-cabinet Cray XC30, called Eos, which was provided by Cray to ensure delivery of the committed INCITE allocations of time in 2013.

- ORNL staff initiated a project exploring the use of digital object identifiers (DOIs), resulting in an invited talk for the 2013 DataCite meeting at the National Academy of Sciences.

COMMUNICATIONS WITH KEY STAKEHOLDERS

Communication with the Program Office

The OLCF regularly communicates with the Advanced Scientific Computing Research (ASCR) Program Office through a series of established events. These include weekly Integrated Project Team calls with the local DOE ORNL Site Office (DOE-OSO) and the Program Office, monthly highlight reports, quarterly reports, the annual Operational Assessment, an annual Budget Deep Dive, and the OLCF annual report. Through a team of communications specialists and writers, the OLCF produces a steady flow of reports and highlights for sponsoring agencies, potential users, and the public.

Communication with the User Community

OLCF communications with users take a wide variety of forms and are tailored to the objective, ranging from relating science results to the larger community or instructing users on more efficiently and effectively using OLCF systems. The OLCF offers many training and educational opportunities throughout the year for both current facility users and the next generation of high-performance computing (HPC) users (see Section 1.4.6).

The impact of OLCF communications is assessed as part of an annual user survey. The mean rating for users' overall satisfaction with OLCF communications was 4.2 in 2013, which was a slight increase from 4.0 in 2012. Eighty-six percent of respondents (303) rated overall satisfaction with communications from the OLCF as "satisfied" or "very satisfied". The OLCF uses various methods to communicate with users, including the following:

- weekly e-mail message
- welcome packet
- general e-mail announcements
- opt-in e-mail notification lists
- Message of the Day
- OLCF website
- conference calls
- OLCF User Council
- one-on-one interactions through liaisons and analysts
- social networking vehicles

Survey respondents indicated that the weekly e-mail message was the most useful form of communication.

Communication with the Vendors

The OLCF conducts formal quarterly reviews of projects and operations with Cray, Inc., and NVIDIA. This process includes specific meetings with the product and program managers, correlation of development schedules across hardware and software products, and field demonstrations of emerging equipment. Early involvement is the key to driving design considerations that positively affect emerging products. Supplementing these formal events, the OLCF meets weekly with its Cray site advocate and Cray hardware and systems analysts to ensure that there is frequent and consistent communication about known issues, bug tracking, and near-term product development.

The OLCF maintains a robust vendor briefing schedule with other product manufacturers as well, making certain that emerging products targeted to this program are well suited to the high-performance, high-capability, and high-capacity needs of the center.

Communication with Advisory Groups

As a part of the Computing and Computational Sciences Directorate (CCSD), the OLCF is subject to a required Directorate Advisory Committee (DAC) process that each year selects a subset of CCSD's divisions, institutes, and projects to review. The purpose of the review is to provide an in-depth assessment of each science and technology directorate, facilitated by the Office of Institutional Planning, to provide assurance that ORNL's science and technology programs and activities meet or exceed performance expectations of sponsors, customers, and major stakeholders. In February 2013, the DAC selected the OLCF as part of their overall review. The DAC was very complimentary of the OLCF's contributions to delivering world-class science, its continued work to address the challenges of exascale, cross-leveraging funding from several sources, and global advocacy for HPC and exascale. The DAC also provided some observations and recommendations to consider as improvement opportunities, which are all being actively pursued.

SUMMARY OF 2013 METRICS

In consultation with the DOE program sponsor, a series of metrics and targets were identified to assess the operational performance of the OLCF in CY2013. The metrics are associated with a series of questions posed to reviewers of the center. The 2013 metrics, target values, and actual results as of December 31, 2013, are summarized below.

The business results metrics are based on the length of time the computational resource has been in production. The results must consider the impact to the calculation of scheduled availability (SA), overall availability (OA), mean time to interrupt (MTTI), mean time to failure (MTTF), and utilization due to "null time" for any system. A null time is a period of time within the reporting period that reduces the total potentially available time. Null times are not considered in the calculation for SA, OA, MTTI, MTTF, and system utilization. The period of time devoted to the Titan acceptance is one example of null time associated with a system.

Summary of the 2013 Metrics

2013 Metric	2013 Target	2013 Actual
Are the processes for supporting the customers, resolving problems, and communicating with key stakeholders and Outreach effective?		
Customer Metric 1: Customer Satisfaction		
Overall score on the OLCF user survey.	Results will be satisfactory (3.5/5.0) based on a statistically meaningful sample.	The OLCF exceeded the metric target: 4.4/5.0.

Summary of the 2013 Metrics

2013 Metric	2013 Target	2013 Actual
Improvement on results that scored below satisfactory in the previous period.	Results will show improvement in at least one-half of questions that scored below satisfactory (3.5) in the previous period.	The OLCF exceeded the metric target: No question scored below satisfactory (3.5/5.0) on the 2013 survey.
Customer Metric 2: Problem Resolution		
OLCF survey results related to problem resolution.	Results will be satisfactory (3.5/5.0) based on a statistically meaningful sample.	The OLCF exceeded the metric target: 4.4/5.0.
OLCF user problem resolution time period.	Eighty percent of OLCF user problems will be addressed within three business days by either resolving the problem or informing the user how the problem will be resolved.	The OLCF exceeded the metric target: 92.3%.
Customer Metric 3: User Support		
Average of user support ratings.	Results will be satisfactory (3.5/5.0) based on a statistically meaningful sample.	The OLCF exceeded the metric target: 4.4/5.0.
Is the facility maximizing the use of its HPC systems and other resources consistent with its mission?		
Business Metric 1: System Availability (For a period of 1 year following a major system upgrade, the targeted scheduled availability is 85% and overall availability is 80%.)		
Scheduled Availability.	Titan: 85% (lower in FY12 due to the compute system upgrades); High-Performance Storage System (HPSS): 95%; external file systems: 95%.	The OLCF exceeded the metric target. Titan: 98.70%; HPSS: 99.99%; Widow0: 100%; Widow1: 99.87%; Widow2: 99.91%; Widow3: 99.87%.
Overall Availability.	Titan: 80%; HPSS 90%; External File Systems 90%.	The OLCF exceeded the metric target: Titan: 93.82%; HPSS: 98.60%; Widow0: 99.11%; Widow1: 97.35%; Widow2: 97.52%; Widow3: 98.09%.
Business Metric 2: Capability Usage		
The OLCF will report on capability usage.	At least 30% of the consumed node-hours will be from jobs requesting 20% or more of the available Opteron nodes.	Capability usage was 59.38%. The OLCF exceeded the metric target.

RESPONSES TO RECOMMENDATIONS FROM THE 2012 OPERATIONAL ASSESSMENT REVIEW

In February 2013 the OLCF presented the 2012 operational activities of the center to the DOE sponsor. Recommendations provided by reviewers, ORNL actions, and DOE ASCR comments and actions are given in the tables below.

1. **Are the processes for supporting the customers, resolving problems, and outreach effective?**
   Assessment: Yes & Yes

Recommendation	ORNL Action/Comments	HQ Action/Comments
None		

2. **Is the OLCF maximizing the use of its resources consistent with its mission? (Financial data will be covered under this question for onsite reviews)**

   Assessment: Yes & Yes

Recommendation	ORNL Action/Comments	HQ Action/Comments
In the 2013 OAR, report progress and results from your work with NVIDIA to get revisions to the NVIDIA driver for more accurate GPU use measurement. Use of the ALTD library tracker, as a proxy for measuring GPU utilization, may not be accurate.	ORNL has worked with NVIDIA and Cray to define requirements for measuring GPU utilization on a per-job basis. NVIDIA remains on schedule for delivering the release-candidate driver in 2Q13. Cray remains on schedule for releasing both CUDA 5.5 and the associated changes to its RUR with CLE 4.2 UP02 in 4Q13. In CY2013, ORNL will evaluate the stability and features of CLE4.2 UP02 and the associated drivers on test and development systems. Given a successful evaluation, ORNL will upgrade to CLE 4.2 on Titan and incorporate RUR in to per-job utilization measurement. This upgrade will then allow these results to be reported in the 2014 and subsequent OARs.	
Once the improved Cray Resource Utilization Software described in Section 2.7.2 is in place, the data gathered should be made available as it would be very useful to the GPU community as a whole.	ORNL will make the per-job utilization information available to the user. This information can supplement data from CrayPat and other performance analysis tools. Without proper context, GPU utilization is not considered appropriate for wide dissemination.	

3. **Is the OLCF enabling scientific achievements consistent with the Department of Energy strategic goals? Specifically applicable to Goal 2: "Maintain a vibrant U.S. effort in science and engineering as a cornerstone of our economic prosperity with clear leadership in strategic areas." Goal 2 includes the targeted outcome: "Continue to develop and deploy high–performance computing hardware and software systems through exascale platforms." Sites may also include contributions to other goals and other targeted outcomes.**

   Assessment: Yes & Yes

Recommendation	ORNL Action/Comments	HQ Action/Comments
None		

4. **Have innovations been implemented that have improved OLCF operations?**

   Assessment: Yes & Yes

Recommendation	ORNL Action/Comments	HQ Action/Comments
None		

5. **Is OLCF effectively managing risk?**

   Assessment: Yes & Yes

Recommendation	ORNL Action/Comments	HQ Action/Comments
None		

6. **Are the performance metrics used for the review year and proposed future years sufficient and reasonable for assessing OLCF's Operational performance?**

   Assessment: Yes & Yes

Recommendation	ORNL Action/Comments	HQ Action/Comments
There is no metric associated with the GPUs on the system. The capability metric is just for CPUs. Given that the bulk of the system performance is encapsulated in the GPUs, OLCF should at least report on GPU usage.	The capability metric is expressed in terms of nodes. Allocation, scheduling, and accounting all consider the contribution of both the CPU and GPU within the node, not just the CPU. That metric is consistent with past descriptions, in that the application must use 20% of the system to be classified as a "capability" job. ORNL will report on the distribution of capability-class jobs over time and by project type. In addition, the OLCF will report on delivered compute hours by GPU-enabled applications.	
While a GPU usage metric is shown for 2013, it is not indicated for 2014 (either as a metric or target). It is recommended that GPU usage be gathered and reported for Titan throughout its lifetime, using a method of calculation agreed to by DOE/ASCR.	Beginning in 2013, consumed hours will denote CPU-only and GPU-enabled contributions to the appropriate reporting categories. We will report this value in 2014 and subsequent years.	
Consider a two-tier capability target such as ALCF has introduced for Mira.	ORNL manages Titan as a single large and homogeneous partition. The capability metric and the accompanying job scheduling policy have ensured that a significant preference is given to larger jobs, while ensuring that overall throughput on the system remains high. ORNL has consistently met or exceeded the capability metric, and does not believe that a two-tiered metric is necessary.	
To DOE HQ: OAR guidance should stipulate the units in which MTTI and MTTF are given. NERSC reports data as days:hours:minutes; ALCF reports data in fractional days; OLCF reports in hours. Either the NERSC or ALCF method is preferred.	OLCF currently reports in fractional hours. As machines get larger and more complex, the MTTI and MTTF metrics are likely to fall below one day so we believe that our reporting unit is appropriate.	

7. What is your overall assessment of the Facility Operational performance?

Recommendation	ORNL Action/Comments	HQ Action/Comments
OLCF continues to execute in an exemplary fashion. Despite significant challenges in deploying Titan, OLCF has been successful in delivering substantial computing resources to their user community, enabling science accomplishments, advancing the frontiers of application performance for GPU computing, and providing other innovations that will be broadly useful within the HPC community.	Thank you.	
The OLCF is well regarded for its effectiveness in managing its user facility. OLCF enables scientific capabilities and advancements that do not exist in the private sector or in academia, and should be commended for the excellent support it provides to its user community. The report provides a thorough description of its planning, risk management strategies, and overall operations. We expect that OLCF will continue to have direct and significant impact on innovation and driving discoveries.	Thank you.	

User Results

HIGH PERFORMANCE COMPUTING FACILITY
2013 OPERATIONAL ASSESSMENT
OAK RIDGE LEADERSHIP COMPUTING FACILITY

February 2014

1. USER RESULTS

CHARGE QUESTION 1: Are the processes for supporting the customers, resolving problems, and outreach effective?

OLCF RESPONSE: Yes. The Oak Ridge Leadership Computing Facility (OLCF) has a dynamic user support model that is based on continuous improvement, regular assessment, and a strong customer focus. One key element of internal assessment is the annual user survey. As part of the survey, users are asked to rate their overall satisfaction with the OLCF on a scale of 1 to 5, with a rating of 5 indicating "very satisfied." The mean rating for overall satisfaction with the OLCF was 4.4, which was an increase from 4.2 in 2012.[*] Overall ratings for the OLCF were positive; 95% of users reported being "satisfied" or "very satisfied." No users reported being "dissatisfied" or "very dissatisfied" with the OLCF overall. The survey also asks users to rate their overall satisfaction with the User Support services provided by the OLCF; the mean rating in 2013 was 4.4.

The center measures its performance using a series of quantifiable metrics. The metric targets are structured to ensure that users are provided prompt and effective support and that the User Support organization responds quickly and effectively to improve its support process for any item that does not meet a minimum satisfactory score. The OLCF exceeded all metric targets for user satisfaction in 2013 with 92.3% of tickets being resolved within three business days. The OLCF continues to enhance its technical support, collaboration, training, outreach, and communication. For example, this year the OLCF broadened the training program to include data processing and analysis. The center also engages in activities to promote HPC to the next generation of researchers.

1.1 USER RESULTS SUMMARY

The OLCF has developed and implemented a user-centric customer support model. The model comprises customer support interfaces, including user satisfaction surveys, formal problem-resolution mechanisms, user assistance analysts, and science, visualization, and data liaisons; multiple channels for communication with users, including the OLCF User Council; and comprehensive training programs, user workshops, and tools to reach and train both current facility users and the next generation of computer and computational scientists. The success of these activities and identification of areas for development are tracked through the annual OLCF user survey.

In an effort to promote continual improvement at the OLCF, users are sent a survey soliciting their feedback regarding support services and their experience as a user of the facility. The 2013 survey was launched on October 2, 2013, and remained open for participation through December 9, 2013. The survey was sent electronically to individuals with active accounts on Innovative and Novel Computational Impact on Theory and Experiment (INCITE), Advanced Scientific Computing Research (ASCR) Leadership Computing Challenge (ALCC), and Director's Discretionary (DD) projects. Three hundred

[*] In this document, "year" refers to the calendar year unless it carries the prefix "FY," indicating the fiscal year.

sixty-seven users completed the survey out of 1,232 possible respondents, with an overall response rate of 30%. This response rate continues to outpace the industry average for closed-end question surveys.

Information was collected about the various users, user experience with the OLCF, and OLCF's support capabilities. Attitudes and opinions on the performance, availability, and possible improvements for the OLCF and its staff were also solicited. Data collected from the user survey were analyzed by the Oak Ridge Institute for Science and Education (ORISE) using both quantitative and qualitative methods. The two fundamental goals that drove the collection and subsequent analysis were to catalog the types of users and to understand their needs. Analysis included basic descriptive statistics and qualitative coding of responses to open-ended questions. Responses to specific survey items were used to cross-check respondents' responses to other items that were directly related to ensure that all responses were valid (e.g., only people who selected that they had used a particular machine could rate their satisfaction with various aspects of that machine). The results of the 2013 survey can be found on the OLCF website.

The effectiveness of the processes for supporting customers, resolving problems, and conducting outreach are defined by the metrics in Table 1.1 and are assessed through the user survey and Request Tracker (RT), the OLCF trouble ticket system.

Table 1.1. 2013 User Result Metrics Summary

2012 Metric	2012 Actual	2013 Metric	2013 Target	2013 Actual
Customer Metric 1: Customer Satisfaction				
Overall score on the OLCF user survey. Target: Results will be satisfactory (3.5/5.0) based on a statistically meaningful sample.	4.2/5.0	Overall OLCF score on the user survey.	Results will be satisfactory (3.5/5.0) based on a statistically meaningful sample.	4.4/5.0
Improvement on results that scored below satisfactory in the previous period. Target: Results will show improvement in at least ½ of questions that scored below satisfactory (3.5) in the previous period.	No question scored below satisfactory (3.5/5.0) on the 2012 survey.	Improvement on results that scored below satisfactory in the previous period.	Results will show improvement in at least one-half of the questions that scored below satisfactory (3.5) in the previous period.	No question scored below satisfactory (3.5/5.0) on the 2013 survey.
Customer Metric 2: Problem Resolution				
OLCF survey results related to problem resolution. Target: Results will be satisfactory (3.5/5.0) based on a statistically meaningful sample.	4.4/5.0	OLCF survey results related to problem resolution.	Results will be satisfactory (3.5/5.0) based on a statistically meaningful sample.	4.4/5.0
OLCF user problem resolution time period. Target: 80% of OLCF user problems will be addressed within three business days, by either resolving the problem or informing the user how the problem will be resolved.	92.3%	OLCF user problem resolution time period.	80% of OLCF user problems will be addressed within three business days by either resolving the problem or informing the user how the problem will be resolved.	92.3%

Table 1.1. 2013 User Result Metrics Summary (continued)

2012 Metric	2012 Actual	2013 Metric	2013 Target	2013 Actual
Customer Metric 3: User Support				
Average of user support ratings. Target: Results will be satisfactory (3.5/5.0) based on a statistically meaningful sample.	4.5/5.0	Average of user support ratings.	Results will be satisfactory (3.5/5.0) based on a statistically meaningful sample.	4.4/5.0

1.2 USER SUPPORT METRICS: USER SATISFACTION

The operational assessment metrics for the HPC facility's user support as assessed by the annual user survey are the following:

- Overall satisfaction rating for the facility is satisfactory.
- Average of user support questions on user surveys is satisfactory.
- Improvement on past year's unsatisfactory ratings occurs as agreed upon with the facility's US Department of Energy (DOE) program manager.

The OLCF metric targets and calendar year (CY) actual results for user support are shown in Table 1.2.

Table 1.2. OLCF User Support Summary: Metric Targets and Calendar Year Results

Survey Area	CY2012		CY2013	
	Target	Actual	Target	Actual
Overall OLCF satisfaction rating	3.5/5.0	4.2/5.0	3.5/5.0	4.4/5.0
Average of user support ratings	3.5/5.0	4.5/5.0	3.5/5.0	4.4/5.0

1.2.1 Overall Satisfaction Rating for the Facility

Users were asked to rate their satisfaction on a 5-point scale, where a score of 5 indicates a rating of very satisfied and a score of 1 indicates a rating of very dissatisfied. The metrics agreed upon by the DOE OLCF program manager define 3.5/5.0 to be satisfactory.

Overall ratings for the OLCF were positive; 95% reported being satisfied or very satisfied with the OLCF overall. With regard to the degree of overall satisfaction with the center, the percentage of satisfied and very satisfied respondents has steadily increased from 2007 (86%) to 2013 (95%).

Key indicators from the survey, including overall satisfaction, are shown in Table 1.3. They are summarized and broken out by program.

Table 1.3. Satisfaction Rates by Program Type for Key Indicators

Indicator	Mean	Program		
		INCITE	ALCC	DD
Overall satisfaction with the OLCF	4.4/5.0	4.4/5.0	4.4/5.0	4.4/5.0
Overall satisfaction with user assistance	4.4/5.0	4.4/5.0	4.3/5.0	4.4/5.0
Overall satisfaction with Titan	4.2/5.0	4.2/5.0	4.1/5.0	4.1/5.0

1.2.2 Average Rating across User Assistance Questions

The calculated mean of all answers to all user assistance questions on the 2013 survey was 4.4/5.0, indicating that the OLCF exceeded the 2013 user support metric target and that users have a high degree of satisfaction with User Support services. This rating was a slight decrease from 4.5 in 2012. The values for two specific user assistance questions decreased slightly from the previous year: (1) the speed of the initial response to queries and (2) the effectiveness of response to account management queries. In looking at the comments in the section for speed of initial response to queries, 8 users out of 247 noted their dissatisfaction with the response time. Among the comments left by the dissatisfied users, one user commented that more people are needed to answer the phones. At present time, the OLCF only has one phone number: calls go to voicemail when the consultant on duty is on the phone with another user. Another user recommended that the OLCF respond more quickly to initial help tickets. In reviewing the comments in the effectiveness of response to account management query section, 11 users out of 256 indicated they were very dissatisfied with this area. In looking at the comments, the major complaint was that it took too long to establish an account. The OLCF has a rigorous account creation process that requires multiple steps to meet the security requirements to operate the OLCF user facility at a moderate level. These additional steps, including export control reviews and identity proofing, can add additional time to the accounts process.

In response to an open-ended question about the best qualities of the OLCF, user assistance was listed as the top choice by 50% of the survey respondents. The following comments are samples from the survey:

“The OLCF is indispensable to the success of my research. Assistance provided by the OLCF staff is very speedy, efficient, and most helpful.”

“The User Support at Oak Ridge is one of the best I’ve ever experienced.”

“OLCF has been extremely helpful this year to support our requests.”

“The knowledgeable staff are very helpful and it is a great advantage to speak with people who understand what you are trying to do.”

Users were also asked an open-ended question soliciting their top suggestions for improvements. The most common response was to improve the queuing policy. Specifically, most of these respondents indicated they would like better queue policies for smaller jobs. Comments include:

“The constraints of having such large node counts to get reasonable throughput is a bit disappointing.”

“Make it simple to get small amounts of SUs for small projects.”

As a leadership computing facility, the OLCF has a mandate that a large portion of Titan's usage come from large capability jobs. Therefore, the OLCF’s queue policy favors large jobs. If a user reports throughput issues, the OLCF engages in a dialogue with the user to determine if steps such as queue exceptions, queue priority, code improvements, or other actions are appropriate.

Evaluation results for individual support functions are summarized in Table 1.4 and are detailed in the following sections.

Table 1.4. OLCF Satisfaction Ratings for CY2013
Based on Answers to the User Support Survey

Total Respondents: 367

Type of Service	Respondents to Each Question	Rating
User Assistance	338	4.4/5.0
Training and Education		
In Person	56	4.3/5.0
Webcast	63	4.2/5.0
Communications	351	4.2/5.0
Website	320	4.1/5.0

1.2.2.1 User Assistance Evaluation

- For support services used, 59% of the 337 respondents reported using User Assistance, followed by 33% using the Scientific Computing/Liaison service.
- Overall satisfaction with the User Support services provided by the OLCF was high with a mean response of 4.4/5.0, which was equal to the rating in the 2012 survey results.
- Respondents with at least one interaction with User Assistance were asked about the speed of initial contact and the quality of the response; a large percentage was satisfied or very satisfied with the initial contact (89%) and with the quality of the response (86%).

1.2.2.2 Training and Education

- The mean rating to the question of overall satisfaction with live in-person OLCF training events was 4.3, which was an increase from 4.0 in 2012. The majority of OLCF users said “yes” (75%) or “maybe” (23%) to the prospect of attending future OLCF training events in person, based on their previous experience.
- The mean rating to the question of overall satisfaction with live OLCF training events via webcast was 4.2. The majority of OLCF users said “yes” (86%) or “maybe” (14%) to the prospect of attending future OLCF training events live via webcast, based on their previous experience.
- The number one reason users gave for not participating in any in-person training events was that they do not have the time to attend.
- When presented with a list of training topics, respondents’ most frequently requested topic was graphics processing unit (GPU) programming (61%) followed by tuning and optimization (49%), and advanced message passing interface (MPI) (46%).
- When asked what the OLCF did well regarding training, 50% noted that the quality of the speakers was what they liked best about the training events. When asked how the OLCF could improve, 29% noted that the OLCF should make the training materials available on line ahead of time.

1.2.2.3 OLCF Communications

- The mean rating for users' overall satisfaction with OLCF communications was 4.2 in 2013, which was an increase from 4.0 in 2012.

- Eighty-six percent of respondents rated their overall satisfaction with communications from the OLCF as satisfied or very satisfied. One user indicated they were very dissatisfied with the communication efforts. The user indicated that the dissatisfaction was due to the large number of survey email reminders they received.

- Respondents indicated that the e-mail message of the week was most useful; Twitter was found to be the least useful communication mechanism.

1.2.2.4 OLCF Website

- The mean rating for users' overall satisfaction with the OLCF website was 4.1 in 2013, which was a slight improvement from 4.0 in 2012.

- Ninety-nine percent of respondents indicated that they had visited the OLCF website. Of these users, 32% indicated that they visit the site once a week or more. Only four respondents indicated they had never visited the site.

- In thematic responses for the OLCF website, it was noted that the OLCF could make improvements to the "My OLCF" dashboard. Several enhancements are already under way and are slated for production in 2014.

1.2.3 Improvement on Past Year Unsatisfactory Ratings

Each year the OLCF works to show improvement in no less than half of any questions that scored below satisfactory (3.5/5.0) in the previous year's survey. All questions scored above 3.5 on both the 2012 and 2013 surveys.

1.2.4 Assessing the Effectiveness of the OLCF User Survey

Before sending the survey, the OLCF met with the ORISE evaluation specialist to review the content of the survey questions to ensure that they accurately addressed the concerns of the OLCF and that all technical terminology was appropriately used. The evaluator specifically reviewed the response options for each of the selection items and discussed how variations in question type could influence the meaning and utility of the data they would generate.

Several targeted notifications were sent to those eligible to participate in the survey. The initial survey invitation from ORISE was sent on October 2, 2013, and subsequent follow-up reminders were sent by Arthur Bland (OLCF project director), Ashley Barker (User Assistance and Outreach group lead), Jack Wells [National Center for Computational Sciences Division (NCCS) director of science], ORISE, and individual members of the OLCF. The survey was advertised on the OLCF website and was mentioned in the weekly communications e-mail sent to all users. Survey responses were tracked on a daily basis to assess the effectiveness of the various communication methods (see Appendix A). The notifications from center management were the most effective, but the results show that other efforts, such as including the notice in the weekly communication, also contributed to the survey response rate.

The distribution of survey respondents was relatively equally balanced in terms of their length of time using the systems (Table 1.5).

Table 1.5. User Survey Participation

	2012 Survey	2013 Survey
Total Number of Respondents (Total Percentage Responding to Survey)	386 (38%)	367 (30%)
New Users (OLCF User < 1 Year)	33%	43%
OLCF User 1–2 Years	25%	26%
OLCF User > 2 Years	42%	31%

Survey respondents were asked to classify the program types with which they were affiliated. Table 1.6 is a summary of responses according to respondent affiliation and program type.

Table 1.6. User Survey Responses by Affiliation and Program Type

Category	Response Rate[a]
Affiliation (*b*)	
University	52%
DOE/Lab/Government	27%
Other	9%
Industry	7%
Foreign	5%
Program Type (*b*)	
INCITE	55%
DD	46%
ALCC	19%
Other	0.3%

* Total is greater than 100% because survey respondents can be associated with more than one type of project.
[b] Percentage of the original survey list.

1.2.4.1 Statistical Analysis of the Results

Statistical analysis of four key survey areas is shown in Table 1.7. They reflect overall satisfaction with the facility, services, and computational resources.

Table 1.7. Statistical Analysis of Key Results

	Overall Satisfaction with the OLCF	Overall Satisfaction with User Assistance	Overall Satisfaction with OLCF Communications	Overall Satisfaction with Titan
Number of Survey Respondents	367	367	367	367
Number of Respondents to This Specific Question	338	338	351	282
Mean	4.4/5.0	4.4/5.0	4.2/5.0	4.2/5.0
Variance*	0.34	0.53	0.47	0.44
Standard Deviation*	0.59	0.73	0.69	0.66

*The OLCF examined the variance and standard deviation for several key questions and found them to be within acceptable parameters.

1.3 PROBLEM RESOLUTION METRICS

The operational assessment metrics for problem resolution are the following:

- Average satisfaction ratings for questions on the user survey related to problem resolution are satisfactory or better.

- At least 80% of user problems are addressed (the problem is resolved or the user is told how the problem will be handled) within three business days.

1.3.1 Problem Resolution Metric Summary

In most instances, the OLCF can resolve a reported problem directly, which includes identification and execution of the necessary corrective actions such that the problem is resolved from the users' perspective. Occasionally the OLCF receives problem reports for which its ability to resolve the root cause of the issue is limited due to factors beyond its control. In such a scenario "addressing the problem" requires that the OLCF has identified and carried out all corrective actions at its disposal for the given situation. For example, if a user reports a suspected bug in a commercial product, prudent measures might be to recreate the issue; open a bug ticket with the product vendor; provide the vendor necessary information about the issue; and then provide a workaround to the user, if possible.

The OLCF uses request tracker software to track queries (i.e., tickets) and to ensure that response goals are met or exceeded. Users may submit queries via e-mail, the online request form, or phone. E-mail is the predominant source of query submittals. In addition, the software collates statistics on tickets issued, turnaround times, etc., allowing the OLCF staff to track patterns and to address anomalous behaviors before they have an adverse effect on the work of many users. The OLCF issued 2,409 tickets in response to user queries for 2013 (Figure 1.1). The center exceeded the problem-resolution metric and responded to 92.3% of the queries within three business days, which was the same as 2012 (Table 1.8).

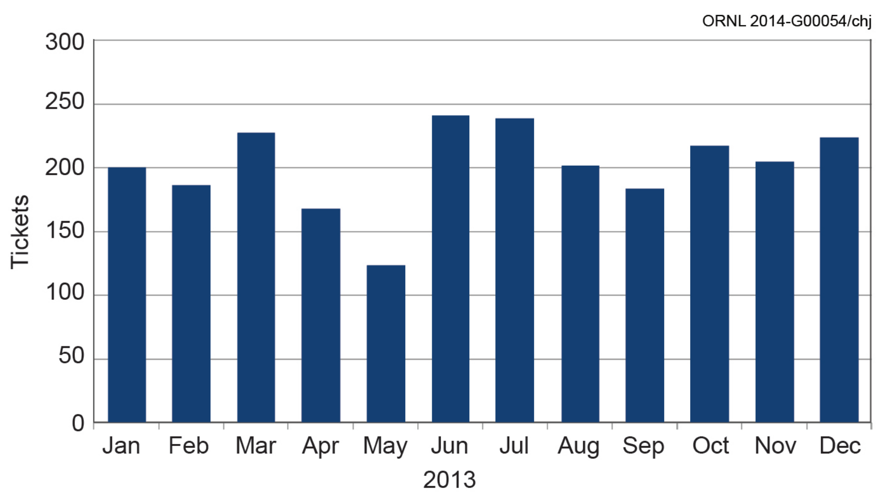

Figure 1.1. Number of Helpdesk Tickets Issued per Month.

Table 1.8. Problem Resolution Metric Summary

Survey Area	CY2012		CY2013	
	Target	Actual	Target	Actual
Percentage of problems addressed in three business days	80%	92.3%	80%	92.3%
Average of problem resolution ratings	3.5/5.0	4.4/5.0	3.5/5.0	4.4/5.0

Each ticket is assigned to one user assistance or account analyst who establishes customer contact and tracks the query from first report to final resolution, providing not just fast service, but also service tailored to each customer's needs. While the OLCF is dedicated to addressing queries promptly, user assistance and account analysts consistently strive to reach the "right" or best solution rather than merely a quick turnaround. Tickets are categorized by their most common types. The top reported problem in 2013 (as well as 2011 and 2012) was related to jobs/batch queues (Figure 1.2).

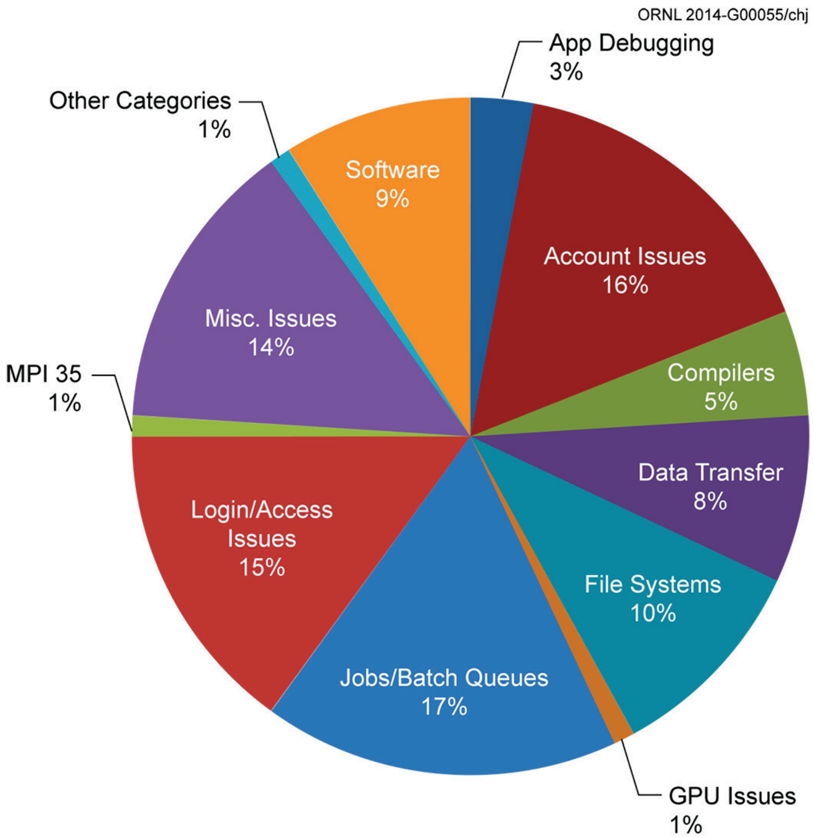

Figure 1.2. Categorization of Helpdesk Tickets.

1.4 USER SUPPORT AND OUTREACH

The operational assessment data for user support and outreach include the following:

- Anecdotal evidence confirms in-depth collaborations between facility staff and the user community.
- A summary of the training events conducted during this period is provided.

The following sections discuss key activities and contributions in the four areas that the OLCF recognizes as the pillars of user support and outreach:

- a user support staff made up of account management liaisons, User Assistance analysts, and Scientific Computing Group (SciComp) liaisons;
- multiple vehicles to communicate with users, sponsors, and vendors;
- training events and materials developed and delivered to current and potential users; and
- the strong outreach component needed to interface with the next generation of HPC users, the external media, and the public.

The OLCF recognizes that users of HPC facilities have a range of needs requiring a range of solutions, from immediate, short-term, "trouble-ticket-oriented" support such as assistance with debugging and optimizing code to more in-depth support requiring total immersion in and collaboration on projects. The center provides User Assistance and SciComp groups, two complementary OLCF user support vehicles. SciComp scientific, visualization, and data liaisons are a unique OLCF response to high-performance scientific computing problems faced by users.

1.4.1 User Assistance Analysts

As already discussed in Sect. 1.3, User Assistance (UA) analysts are responsible for addressing user queries. Some of the most common UA activities include the following:

- enabling access to OLCF resources;
- helping users compile and debug large science and engineering applications;
- identifying and resolving system-level bugs in conjunction with other technical staff and vendors;
- installing third-party applications and providing documentation for usage;
- engaging other OLCF staff to ensure that users have up-to-date information about OLCF resources and to solicit feedback;
- researching, developing, and maintaining reference and training materials for users;
- communicating with users;
- developing and delivering training; and
- acting as user advocates.

1.4.2 SciComp Liaisons

A team of experts in a broad range of computational, visualization, data, and computer science disciplines and with extensive experience using the center's computer resources are available to support user project teams. Science liaisons assist with algorithm and performance improvements of the science codes. Visualization liaisons design and develop the visualization and analytics capabilities needed to support the data interpretation and presentation requirements. Data liaisons assist users with innovative solutions for input/output (I/O) functionality in scientific applications and large-scale data workflow challenges. See the science highlights in Section 3.2.1 and 3.2.2 for examples of data liaison support. For many of the INCITE projects, the support role of each of the liaisons is better described as a collaborative partnership with the science teams. Of the 56 SciComp publications in 2013, 25 were directly related to INCITE and OLCF responsibilities.

In the first half of 2013 SciComp was led on an interim basis by Bronson Messer. In May 2013, Tjerk Straatsma took on the role SciComp group leader. Judy Hill, the SciComp science liaison task lead, works with the science teams to develop a long-term strategy for the liaison support roles within the OLCF.

Liaisons carry out the computational readiness review of proposals, may participate as observers of the scientific review panels, and contribute input to the final selection meeting, where INCITE projects are recommended for allocations. Each awarded INCITE project is assigned one or more liaisons, taking care to match the expertise of the liaison to the project.

SciComp includes postdoctoral fellows as integral participants in the OLCF mission, enabling additional science to be accomplished as well as training future computational domain scientists. The science liaison task lead manages the ARRA Post-Doctoral Program, which employed six postdoctoral fellows during 2013 and will end August 31, 2014. The responsibilities include ensuring that postdoctoral fellows receive appropriate guidance and oversight from their technical mentors and that they have appropriate financial and computational resources.

An example of a key effort by an OLCF postdoctoral fellow in 2013 was the development of a locally self-consistent multiple scattering code (LSMS) driver that implements the framework of the recently proposed replica-exchange parallel Wang-Landau (WL) scheme. This is the work of postdoctoral associate Ying Wai Li. She is examining different mechanisms for code improvement, in particular to address the effect of I/O on the scalability and performance of WL-LSMS [Vogel et al. (forthcoming), Vogel et al. 2013].

Related Publications:

T. Vogel, Y. W. Li, T. Wüst, and D. P. Landau, "Exploring new frontiers in statistical physics with a new, parallel Wang–Landau framework," *J. Phys.: Conf. Ser*. (forthcoming).

T. Vogel, Y. W. Li, T. Wüst , D.P. Landau, "A generic, hierarchical framework for massively parallel Wang-Landau sampling," *Phys. Rev. Lett*. **110**, 210603 (2013).

1.4.3 Center for Accelerated Application Readiness

The Center for Accelerated Application Readiness (CAAR), established in 2009, carried out performance improvement activities after the upgrade from the Fermi GPUs to the Kepler GPUs for the scientific applications listed in Table 1.9. Codes such as WL-LSMS have been shown to effectively scale across the entire machine. All of the codes listed are being used by large-scale projects on Titan in the INCITE or ALCC programs. Accomplishments made possible as a result of the CAAR efforts include the use of the LAMMPS molecular modeling code in the "Film Dewetting" and "Organovoltaic Materials" Early Science projects highlighted in Sections 1.4.4 and 3.2.4. Another example is the materials science application WL-LSMS that not only has enabled new scientific insights into magnetic materials as part of Markus Eisenbach's INCITE project, "Scalable First Principles Calculations for Materials at Finite Temperature," but also has demonstrated that the use of GPU acceleration leads to a 8.6-fold reduction in runtime accompanied by a measured reduction in consumed energy with a factor of 7.3 for runs that have used 18,561 Titan nodes.

Table 1.9. Application Performance Benchmarking

Application	Performance Ratio		
	XK6 (Fermi GPU) vs XK6 (w/o GPU)	**XK6 (Fermi GPU) vs XE6**	**XK7 (Kepler GPU) vs XE6**
S3D	1.5	1.4	2.2
Denovo	3.5	3.3	3.8
LAMMPS	6.5	3.2	7.4
WS-LSMS	3.1	1.6	3.8
CAM-SE	2.6	1.5	*a*

[a] Performance tuning for CAM-SE on the XK7 has not yet been completed.

1.4.4 Scientific Support Examples

Examples reported here are representative of the support provided to research campaigns awarded time through the INCITE, ALCC, and DD programs. Also included are examples of Early Science support provided by postdoctoral associates. The Early Science (ES) program is a resource allocation program on the Titan supercomputer in which application teams associated with CAAR were provided early access to Titan for the purpose of achieving specific science goals as well as to demonstrate the impact of their code-porting activities.

1.4.4.1 Modeling Core Collapse Supernovae

Bronson Messer of SciComp collaborated with users of Anthony Mezzacappa's INCITE project, "Three Dimensional Simulations for Core Collapse Supernovae," to improve the scalability and physical fidelity of the CHIMERA code. His contribution to the software development consisted of adding OpenMP threading to both the neutrino transport and nuclear burning modules in the code and resulted in substantial speedups (Messer et al. 2013). These performance improvements have allowed the collaboration to increase the number of species in the nuclear burning modeled with CHIMERA from 14 to 150 isotopes at exactly the same throughput rate. This kind of increased fidelity will allow explosion simulations to make quantitative, testable predictions of nucleosynthesis—the formation of nuclear species in a supernova event and the primary source of all elements heavier than iron in the periodic table. Work is continuing on porting the nuclear kinetics modules to Titan's GPU accelerators.

Related Publication:

O. E. B. Messer, J. A. Harris, S. T. Parete-Koon, and M. A. Chertkow, "Multicore and accelerator development for a leadership-class stellar astrophysics code," Lecture *Notes in Computer Science* 7782, **92** (2013).

1.4.4.2 The Nuclear Quantum Many-Body Problem

Hai Ah Nam is a member of the SciDAC-3 project "Nuclear Computational Low-Energy Initiative," a large-scale collaborative team designing highly optimized modeling capabilities for the development, validation, and application of a novel nuclear energy density functional. The computational needs for this effort are supported by James Vary's INCITE project, "Nuclear Structure and Nuclear Reactions," for which Nam is the OLCF liaison. This project is focused on using the theory of strong interactions and quantum chromodynamics to develop a unified description of nuclei and accurate predictions of nuclear reactions. A recent publication in *Computer Physics Communications* (Bogner et al. 2013) gives an extensive review of the algorithmic and computational implementation developments from this collaborative effort.

Related Publication:

S. Bogner, A. Bulgac, J. Carlson, J. Engel, G. Fann, R. J.Furnstahl, S. Gandolfi, G. Hagen, M. Horoi, C. Johnson, M. Kortelainen, E. Lusk, P. Maris, H. Nam, P. Navratil, W. Nazarewicz, E. Ng, G. P. A. Nobre, E. Ormand, T. Papenbrock, J. Pei, S. C. Pieper, S. Quaglioni, K. J. Roche, J. Sarich, N. Schunck, M. Sosonkina, J. Terasaki, I. Thompson, J. P. Vary, and S. M. Wild, "Computational Nuclear Quantum Many-Body Problem: The UNEDF project," *Computer Physics Communications* **18**4, 2235–2250 (2013).

1.4.4.3 Turbulent Combustion

OLCF's Ramanan Sankaran developed a suite of highly optimized GPU-enabled computational unsteady flamelet routines for turbulent reactive flow applications (Sankaran 2013). The applications can use the suite of routines within codes that use conventional processors to carry out the legacy flow solvers. The implementation of the flamelet models exposes parallelism and concurrence at multiple levels, which enable the effective utilization of GPU capabilities. The improved device-scale simulation capabilities for practical engineering applications contribute to a better understanding of the fundamental interactions between turbulent flow and the chemical reactions in combustion.

Related Publication:

R. Sankaran, "GPU-accelerated software library for unsteady flamelet modeling of turbulent combustion with complex chemical kinetics," *51st AIAA Aerospace Sciences Meeting including the New Horizons Forum and Aerospace Exposition, American Institute of Aeronautics and Astronautics,* doi: 10.2514/6.2013-372 (2013).

1.4.4.4 Thin Film Dewetting: An Early Science Research Project

In support of an OLCF Early Science project, Postdoctoral associate Trung Ngyen is investigating the application of efficient methods for long-range electrostatic interactions in massively parallel molecular simulations using the Large-Scale Atomic/Molecular Massively Parallel Simulator (LAMMPS) on Titan (Ngyuen et al., in press). His study of the effects of molecular ordering and film thickness on the process of thin film dewetting is the largest reported to date. The large-scale simulations avoid finite-sized effects that might suppress the growth of unstable capillary modes and suggest a common origin for spinodal instabilities and thermal nucleation. These studies have provided new insights into the rupture events that occur in liquid-vapor equilibrium.

Related Publication:

T. D. Ngyuen, J-M. Y. Carrillo, M. A. Matheson, and W. M. Brown, "Rupture mechanism of liquid crystal thin films realized by large-scale molecular simulations," *Nanoscale* (in press).

1.4.4.5 Organic Photovoltaic Materials Simulations: an Early Science Research Project

Postdoctoral associate Jan-Michael Carrillo is studying the behavior and interactions of polymers at surfaces and interfaces, including interactions of charged or neutral polymers to nanoparticles and biomaterials, using coarse-grain molecular dynamics simulations. His current work focuses on the effects of substrates and additives on the morphology of the bulk heterojunction in organic photovoltaic materials (Carillo et al. 2013). In collaboration with researchers at the Center for Nanophase Materials Sciences at ORNL and Jülich Centre for Neutron Science, for this Early Science project he is modeling and simulating neutron spin echo experiments, specifically the dynamics of confined flexible polymers, and is comparing them to experimental results.

Related Publication:

J-M. Y. Carrillo, R. Kumar, M. Goswami, B. G. Sumpter, and W. M. Brown, "New insights into the dynamics and morphology of P3HT:PCBM active layers in bulk heterojunctions," *Physical Chemistry Chemical Physics* **15** (41), 17873–17882 (2013).

1.4.4.6 Visualization

The visualization liaisons provide support for an impressive number INCITE, ALCC, DD, Early Science, Industrial HPC Partnership Program, and other users. This benefits the respective projects directly, but it also highlights the integrated visualization capabilities provided by the OLCF and illustrates the level of integrated expert support that the center provides to help the domain science teams be productive and successful. Among the many examples of visualization contributions is the illustration of how atomic motions in cellulases are responsible for the binding to cellulose and the subsequent enzymatic processes important for bioenergy production, which is the key scientific challenge addressed in Jeremy Smith's INCITE project, "Cellulosic Ethanol: Simulation of Multicomponent Biomass Systems."

1.4.5 Outreach and Communications

The OLCF works to engage new and next-generation users and showcases research at the OLCF through strategic communication activities such as tours, highlights, fact sheets, posters, snapshots, the OLCF website, and center publications. The OLCF was responsible for the creation of 65 highlights and for more than 140 total outreach products in 2013 (see Appendix B). Throughout the year, the OLCF provides tours to groups of visitors who range from middle-school students through senior-level government officials. The center gave tours for 281 groups in 2013.

The OLCF website received 241,038 visits and 410,999 page views in 2013. The graphic design community honored the OLCF web designer, Brian Gajus, in 2013. Gajus took an American Inhouse Design Award for his work on the OLCF Titan page, which introduced America's most powerful supercomputer. Sponsored by the news magazine *Graphic Design USA*, the award recognizes the contributions made by designers working in institutional marketing and communications departments. Gajus also won a MarCom Gold Award for the same site. MarCom Awards is a creative competition for any individual or company involved in the concept, writing, and design of print, visual, audio and web materials and programs. Entries come from corporate marketing and communication departments, advertising agencies, public-relations firms, design shops, production companies, and freelancers. With 119,925 page views, the Titan page was among the most visited sites on the OLCF website.

The OLCF science writers and graphic designer were also MarCom award recipients. The team of science and writers and the graphic designer were recognized with a Platinum Award for the 2013 OLCF annual report. The team also won a MarCom Gold Award for the story, "Titan's New Build Attracts Magnetic Systems Research Impossible Until Now."

OLCF science writers produce highlights on a regular basis on topics ranging from science results to technological advances to professional achievements. These articles are often released externally through a "ping" to appropriate media. Often, those highlights are picked up for external publication. One example of this was an article titled "Researchers recruit Titans to study key molecular switch that controls behavior." The science results were published in *Nature Communications*, and the resulting highlight was released to the media. The story was picked up by 20 different media outlets, including: *Scientific Computing*, *R & D Magazine*, *Technology News*, and Press News.org to name a few.

Highlights, such as "Extermination at Scale," (Allinea News), "Titan Simulates Earthquake Physics Necessary for Safer Building Design" (EurekaAlert) and "Titan Speed Ups" (HPCWire.com) were picked up for external publication. Other examples include: "Vampir Rises to the Occasion at ORNL" (HPCWire.com), "Big Rig Design Still Going Strong" (HPCWire.com), and "INCITE Program Doles Out Supercomputer Hours" (Energy.gov). The group continues to track media usage of these highlights as a part of its ongoing effort to maximize the breadth and depth of news coverage of the OLCF.

The OLCF partnered with CSCS, the Swiss National Supercomputing Centre, and the ALCF in leading Birds of a Feather (BoF) meetings at the ISC2013 and SC2013 conferences on the theme of "High-Performance Communication for High-Performance Computing." These BoF sessions were well attended, bringing together leaders from supercomputing centers, universities, industry, and associated fields (e.g., science journalists and HPC solution providers) to discuss communicating the value of supercomputing for society at large. Outcomes of these meetings include the creation of the hpc-hpc web portal and email distribution to facilitate sharing of information.

Late in the year, the outreach team worked with the editor of a French-based publication, *HPC Magazine*, to set up a comprehensive interview for the NCCS director of science, focusing on Titan and its vast capabilities. Further, the team led a multidiscipline group of professionals who helped the director of science prepare for the interview. It is just this type of cross-functional effort that typifies the type of outreach communications in which the team engages on an ongoing basis.

1.4.5.1 Titan User Guide

As reported in the 2012 OAR, the OLCF worked to implement new code within its website content management system to present a curated subset of articles as a single system user guide for an individual OLCF system. This technique leverages existing content (thereby reducing duplication of effort) and retains the ability to search for the individual articles for those users who prefer that approach while enabling the creation of a single definitive guide with a logical flow for those users who prefer a narrative structure. In 2012, the OLCF worked to populate the *Titan User Guide*[†] in preparation for system rollout, and more work was done in 2013 to update and add to the documentation. The *Titan User Guide* was the third most visited page on the OLCF website in 2013, with over 7,657 page visits from 4,946 unique visitors, indicating that the guide is being used by more than just the OLCF user population. OLCF users were asked to rate their satisfaction with the guide on the 2013 survey. It received a rating of 4.2, indicating that users are satisfied with the work done to assemble the guide. As the OLCF gained more experience with Titan, the staff continued to add new documentation to the guide. A few comments from the survey include:

> "User guides are very helpful on the website "
>
> "Overall I found the documentation pretty easy to find and to use. Up to this point I have been able to find everything I needed in the online docs. My suggestion is merely to keep doing what you're doing, as I've generally been pleased with the online documentation for Titan"

1.4.5.2 Online Tutorials

The OLCF continued to develop new tutorials for the OLCF website, partly in response to feedback received in prior surveys. The tutorials section of the OLCF website received 12,325 page views in 2013, with 10,715 unique visits. The CUDA vector tutorial added in 2012 was the ninth most visited page on the entire OLCF website in 2013 with 2,260 page views. The OLCF added nine new programming tutorials in 2013:

- Accelerating Serial Code for GPUs
- Accelerator Interoperability
- Accelerator Interoperability II
- Compiling Mixed GPU and CPU Code
- CUDA Proxy: Managing GPU Context
- GPUDirect: CUDA aware MPI
- OpenACC Game of Life
- PETSc
- Serial to Parallel: Monte Carlo Operation

Users were asked to rate, for the first time, the OLCF tutorials found on the OLCF Tutorials web page. Users indicated they were satisfied with the tutorials, with a mean rating of 4.0. Comments from the survey include the following:

> "Please keep the online tutorials coming! These have been incredibly useful."
>
> "I am still a rookie at this stuff, so your web page tutorials are of critical importance to me."

[†] https://www.olcf.ornl.gov/support/system-user-guides/titan-user-guide/

1.4.5.3 Monthly User Calls

Due to the very low participation in the monthly user conference calls, in 2011 the OLCF moved to a quarterly format. In an effort to reach users more frequently and outside the formality of training events, monthly user calls were reinstated in 2013 but employed a new approach: a short (~10 min) live presentation on a topic of interest called a "nugget," after which users are invited to speak to members of the OLCF staff. These nuggets have covered a wide range of topics, such as data transfer methods, introduction of new resources, usage of new tools, and large policy changes (e.g., file system migration, queues). Attendance to these calls has steadily increased, and feedback from users indicates that they have been valuable. After one of the calls regarding data transfer, a user emailed to say:

> "I wanted to thank you both for the presentation last week about Globus online and grid transfers, and for persisting with the call despite the technical difficulties. I thought I'd send you some feedback on this.
>
> I've managed to get a new OSG cert (rather than replacement) like Hai Ah suggested. With help from Mitchell Griffith, I got it registered and set up and got the JLab endpoint registered (you've probably caught some of the helpdesk traffic on that). Since the Tuesday Call, I have brought over the entire production for this and last year's INCITE which is great.
>
> So I just wanted to say 'Thanks!' and also to pass on my very best wishes and thanks to your teams who set this up, and produced the nice documentation and persisted with the call."

1.4.6 Training

Workshops, user conference calls, training events, and seminars are integral components of the OLCF user support model. While training can obviate difficulties in doing science on such large-scale systems, training events can also serve to engage both the public and the user community. In 2013, the OLCF training program focused on increasing overall engagement with users and the greater community by creating more opportunities for participation. The OLCF hosted 30 events, with more than 750 participants. With Titan coming on line, user training focused on programming and using accelerators.

The OLCF collaborated with other HPC facilities with the goals of increasing efficiency by sharing resources, increasing the quality of the training material by involving more experts, and increasing user participation in training events. The OLCF collaborated with the National Institute for Computational Sciences to jointly teach a training class on the new XC30 architecture. By combining resources and expertise, the organizers improved the quality of the material and shared resources needed to cover the training class. The OLCF also collaborated with the National Energy Research Scientific Computing Center (NERSC) to advertise their training events to OLCF users and vice versa. NERSC held an event titled "Introduction to High Performance Computing Using GPUs" in July. The OLCF advertised this event to OLCF users and hosted a WebEx session locally for ORNL users to participate in the workshop. The OLCF also worked with NERSC to advertise a few of its training events and NERSC staff was gracious enough to participate and send a speaker to the data workshop. The OLCF also participated in the Argonne Training Program for Extreme-Scale Computing (ATPESC) event by issuing user accounts on Titan, participating in some of the sessions, and answering questions about using Titan. Members from the OLCF, NERSC, and the ALCF have formed a training collaboration group that has begun meeting regularly with the goal of hosting joint training events and sharing training resources.

See Table 1.10 for a partial list of training events. See Appendix C for the complete list of 2013 training events.

Table 1.10. Training Event Summary

Event Type	Description	Date	Participants
Workshop/Training	OOP (Object-Oriented Programming) Workshop	Feb 5-7, 2013	68
Seminar Series	Filippo Spigaj, University of Cambridge, “Quantum ESPRESSO”	March 15, 2013	Unknown
Seminar Series	Joost VandeVondele, ETH Zurich, “Density Functional Theory (DFT) Based Simulation”	April 5, 2013	Unknown
Workshop/Training	LQCD Workshop	April 29–May 3, 2013	25
Seminar Series	Thomas Wuest, Swiss Federal Research Institute WSL, “Unknotting Challenging Questions in Protein Physics using Wang-Landau Sampling”	May 23, 2013	Unknown
Workshop/Training	Programming with Big Data in R: pbdR	June 17, 2013	22
Seminar Series	Karol Kowalski, Pacific Northwest National Laboratory, “The Coupled Cluster Formalism Across Spatial, Energy, and Time (to solution) Scales”	June 27, 2013	Unknown
Workshop/Training	GPU Programming for Molecular Modeling Workshop	August 3, 2013	12
Workshop/Training	Processing and Analysis of Very Large Data Sets	August 6-8, 2013	80
Seminar Series	Michael S. Pindzola, Auburn University, “Atomic and Molecular Collisions using a Time-Dependent Close-Coupling Method”	August 16, 2013	Unknown
Workshop/Training	ADIOS Code Spring Workshop	August 19-20, 2013	19
Seminar Series	Stan Tomov, The University of Tennessee, “High-Performance Linear Algebra with Intel Xeon Phi (MIC) Coprocessors”	September 9, 2013	Unknown
Seminar Series	Patrick Charbonneau, Duke University, “High-Dimensional Surprises Near the Glass and the Jamming Transitions”	October 28, 2013	Unknown
Seminar Series	Rebecca Hartman-Baker, iVEC, “Enigmas in the Outback: Computational Science at Unprecedented Scales”	December 3, 2013	Unknown

1.4.6.1 East/West Coast Titan Workshops

In attempt to minimize travel costs for the users, the OLCF conducted two similar training events in January and February of 2013 in opposite geographic locations. The format was intended to help facilitate access to training by users. It also served to increase collaboration with the OLCF accelerator vendor/partner NVIDIA, who graciously offered a classroom for the first event within its headquarters in San José, California. The second was held in downtown Knoxville, Tennessee, at a Hilton, in response to user feedback that indicated difficulties in reaching the ORNL campus and additional transportation costs. Both events had an almost identical curriculum, which included hands-on exercises on GPU programming, practicums on tools, and best practices learned by OLCF’s staff and the CAAR team. As with all of our training events, both were broadcast live over the web and reached more than 176 people both remotely and on site.

1.4.6.2 Processing and Analysis of Very Large Data Sets

In 2013 the OLCF broadened the training program to include data processing and analysis. It became apparent upon completion of the *Application Requirements for Exascale* report[‡] that users were struggling with analyzing increased data sets being generated by Titan. The workshop lasted three days and covered major aspects of data processing such as I/O, scalable data tools, and visualization. In addition to ORNL staff, the workshop also included speakers from Brookhaven National Laboratory (BNL), Lawrence Berkeley National Laboratory (LBNL), the University of Chicago, the Swiss National Supercomputing Centre, the University of Tennessee, and Kitware.

1.4.6.3 Training the Next Generation

The OLCF maintains a broad program of collaborations, internships, and fellowships for young researchers. Thirty-one faculty, student interns, and postdoctoral researchers were supported fin 2013. Examples of user engagement and outreach include the following:

- Benjamin Brock, a University of Tennessee (UT) Haslam Scholar and computer science major, served as an intern from May until August. Under the supervision of OLCF research scientist Judith Hill, Brock helped port an out-of-core algorithm for solving systems of equations known as LU Factorization to UT's Beacon supercomputer. Dense LU Factorization is well understood and widely established as a supercomputer benchmark to rank the performance of systems for the TOP500 list. "By porting an LU Factorization algorithm, I can evaluate the performance of Beacon's Intel Xeon Phi coprocessors and investigate the development tools and libraries that are available, comparing them to the analogous GPU development tools and libraries on Titan," Brock said.

- Pellissippi State Community College student Jake Wynne III interned part-time at the OLCF from May until December 2013. Under the supervision of Suzanne Parete-Koon, an OLCF user support specialist, Jake wrote tutorials that demonstrate parallel application programming interfaces (APIs) such as MPI, openMP, and CUDA. Wynne developed a step-by-step process for writing tutorials that show the progression from a completely serial code, meaning no APIs, first to MPI then to openMP, and finally to CUDA. "These tutorials are vital for Titan users to know how to get the most out of parallelization," he said.

- The youngest intern in ORNL's history was 13-year-old William Walker Smith, who participated in the Science Saturdays program. This was the first year for Science Saturdays, which is sponsored by ORNL and administered by Oak Ridge Associated Universities (ORAU). The 10-weekend program allowed students in grades 8 through 12 to attend ORNL scientists' lectures and to participate in hands-on activities beginning in February. Smith landed an internship working with OLCF SciComp member, Hai Ah Nam, by being one of 10 winners in a related essay contest. The upcoming eighth-grader's essay was adapted from the research he did for the Southern Appalachian Science and Engineering Fair (SASEF). "At SASEF, I tested three methods for gathering solar energy," says Smith. "Using that research, I proposed a method for making photovoltaic cells more efficient." This summer, Smith turned his attention from solar energy to programming in Java. In a talk he gave to OLCF staff and researchers, he explained that nontraditional education tools for programming, such as Khan Academy and Code Spells, could become part of school curricula as early as middle school.

[‡] *Scientific Application Requirements for Leadership Computing at the Exascale*, ORNL/TM-2007/238, Oak Ridge National Laboratory, Oak Ridge Tennessee.
https://www.olcf.ornl.gov/wp-content/uploads/2010/03/Exascale_Reqms.pdf

- Philip Curtis, a senior in the computer science program at Tennessee Tech in Cookeville, and Chris Martin, an Oak Ridge High School graduate, worked closely with the OLCF's Jim Rogers to give insight into some of Titan's operations. Martin created a 3D view of Titan that pulls real-time sensor data from each cabinet and node, including CPU, GPU, and DIMM temperatures, and allows a person to select different statistical views of that data, including minimum, maximum, and average values, to get a better understanding of how Titan is performing. Although Curtis worked on multiple projects, one way he gave insight into Titan's operations was by creating an interface that looks first at CPU power utilization data for an application, and then CPU and GPU power utilization data for the GPU-enabled version of that same application. "By looking at some of these operations, for example power consumption, we can defend Titan's use of GPUs," said Curtis. "GPUs require more energy to run, but the time frame to complete the run is significantly less, making Titan more cost-effective in the long run."

- During SC13, OLCF staff member Fernanda Foertter led a "birds-of-a-feather" session speaking on women in computing and stressing the need for better recruiting and retention policies for women working in science, technology, engineering, and mathematics.

- For the past 4 years Dustin Leverman has participated in the Student Cluster Competition as a Supercomputing (SC) committee member; in SC13 Dustin was the Cluster Competition committee chair. The competition exposes students to HPC in a competitive environment, where they race to build the fastest application under significant energy constraints. By participating in this event, the OLCF has the opportunity to engage with students interested in HPC and as a result has hired two former competition participants, including Dustin Leverman himself.

- In an effort to reach out to "future users," the training program added classroom support to professors covering aspects of scientific computing and HPC. This support includes, but isn't limited to, on-site lectures given by UA staff, access to training accounts, exercises, and curricula. Two early adopters included a Mathematical Methods course taught by Professor Ken Read at the University of Tennessee, and a Distributed Computing course taught by Professor Chris Lupo at California Polytechnic. More classroom visits are planned. This outreach is also an opportunity to instruct faculty about the availability of center resources and describe the skills required to work in HPC. The program also serves as a way to establish a relationship with faculty for referrals to OLCF pre-professional programs (e.g., internships, postdocs).

1.4.7 Other Notable Support Activities

In addition to those highlighted above, other sample activities include the following:

- The UA accounts team was responsible for the maintenance of 2,320 unique OLCF accounts and 254 projects for users, staff, and vendors in 2013. The accounts team also participated in the DOE Office of Inspector General (IG) Audit of User Facilities in 2013 with the end result of no findings issued. This team worked diligently to provide a tremendous amount of data to the IG over a 5-month period.

- The OLCF worked to prepare for the addition of two computing resources, Eos and Rhea, as well as the new Spider II file system. User Support activities included testing, preparing the documentation, communicating with the users, and working to set up the infrastructure needed for job accounting and resource access. The OLCF developed new user guides for Eos and Rhea and created a new *Accelerated Computing User Guide*.

Business Results

HIGH PERFORMANCE COMPUTING FACILITY
2013 OPERATIONAL ASSESSMENT
OAK RIDGE LEADERSHIP COMPUTING FACILITY

February 2014

2. BUSINESS RESULTS

CHARGE QUESTION 2: Is the facility maximizing the use of its HPC systems and other resources consistent with its mission?

OLCF RESPONSE: Yes. The OLCF provides a series of highly capable and reliable systems for the user community. The 2013 reporting period includes the introduction to production of the newest HPC resources, the Cray XK7 Titan, a Cray XC30 (Eos), and the Spider II Lustre file systems. The effectiveness with which these resources were delivered is demonstrated by the business result metrics, which were met or exceeded in all cases. In the case of Titan, OLCF expertise was integral to diagnosing and satisfactorily resolving hardware events that occurred during the initial months of installation and use. The OLCF team communicated with users throughout this period and established policies and job-scheduling priorities that maximized access to the production systems. At year's end, the OLCF delivered all of the compute hours committed to the three major allocation programs: INCITE, ALCC, and DD. OLCF leadership computational resources support scientific research through production simulation across many scientific domains, providing the key computing and data resources that are critical to their success.

2.1 BUSINESS RESULTS SUMMARY

Business results measure the performance of the OLCF against a series of operational parameters. The two operational metrics relevant to the OLCF's business results are resource availability and the capability utilization of the HPC resources. The OLCF additionally describes resource utilization as a reported number, not a metric.

2.2 TITAN—CRAY XK7 RESOURCE SUMMARY

In the fourth quarter of 2011, the OLCF initiated an upgrade of the Cray Jaguar compute system from a model XT5 to a model XK.

- *Phase 1* of this upgrade (2011) included the installation of 4,672 new XK compute blades, each configured with four AMD Opteron™ 6274 processors (18,688 compute nodes total) and the upgrade of the system's interconnect fabric from SeaStar to Gemini.

- *Phase 2* of this upgrade, conducted during the fourth quarter of 2012, included the installation of 18,688 NVIDIA K20X (Kepler) accelerators, where each existing AMD Opteron was connected to the NVIDIA Kepler as a CPU-GPU pair. The completed system, a Cray XK7 with more than 27 petaflops of peak computational capacity is named Titan.

The initial hardware and software work on Titan was completed in October 2012. At that time, the OLCF initiated the comprehensive acceptance test, with explicit hardware diagnostics, functionality tests, performance tests, and stability test elements.

2.2.1 XK7 Compute Blade—PCIe-3 Connector Repair

In December 2012, during the stability test for Titan, the OLCF, Cray, and NVIDIA identified an issue with the PCIe-3 connector that connected each of the NVIDIA SXM GPU cards to the Cray XK7 blade. On a Cray XK7 compute blade, which contains four CPU-GPU node pairs, there is a PCIe-3 connector that provides 16 "lanes" between an individual GPU and the XK7 blade. Each lane contains one signaling pair for receiving data and a second signaling pair for transmitting data. During application execution, the OLCF identified instances of PCIe "lane degrades," where the electronic signaling for some number of the PCIe lanes was lost, affecting communication performance. Root cause analysis indicated that a manufacturing defect related to the soldering process produced a loss of ductility in the solder in some connections. Intermittent failures due to that manufacturing defect were expected to be persistent at higher than acceptable Failure In Time (FIT) rates throughout the anticipated lifetime of the system. Based on that information, the OLCF and Cray initiated a hardware maintenance activity whereby every PCIe-3 connector was removed from service and repaired. This activity was accomplished in two stages, where 96 compute cabinets were moved to a maintenance partition and repaired, and then the other 104 compute cabinets were moved to a maintenance partition and repaired. The period of performance for this maintenance activity was from February 2 to April 11, 2013. This strategy allowed users to continue work on a significant resource while Cray could work an uninterrupted maintenance schedule. At the conclusion of the connector repair work, the entire 200-cabinet system reentered the OLCF acceptance test.

2.2.2 XK7 Acceptance Test

The 200-cabinet Titan partition began the OLCF acceptance test on April 12, 2013, and, over the course of 6 weeks, Cray and the OLCF executed the hardware diagnostics, functionality, performance and stability tests, completing all requirements on May 30, 2013.

The hardware diagnostics verified correct operation of individual components and subsystems, including processors, memory, and the Gemini interconnect. The functionality component of the acceptance test guaranteed correct and complete implementation of the software stack, and verification that the test's scientific applications met the prescribed correctness criteria. The performance component verified the performance and scalability of the interconnect, the file system, and a series of scientific applications over the full scale of the system. The scientific applications included both CPU-only and GPU-enabled code bases. The stability component of the acceptance test verified that Titan could sustain the mix of expected user code development and application workloads. Execution of the job correctness and performance tests was automated through the use of a special test harness, developed by OLCF staff, that managed the scheduling, job failure rates, and correctness criteria measurements and reporting.

2.2.3 XK7 Compute Blade—SXM Mechanical Assembly Revision

With the PCIe-3 connector repair work complete and the system in full production, Cray identified SXM connector-pin FIT rates that were within acceptable margins, but higher than expected. Mechanical stress due to the original SXM mounting scheme was determined to be a fundamental contributing factor. Cray stress-tested several different SXM mechanical assembly mounting changes (design changes), subjecting the mounting changes to repeated thermal-cycle testing and residual testing that were calculated to simulate an entire life cycle of the blade. Through this process, Cray identified a design change that would fully eliminate the stress condition. Beginning September 3, 2013, the OLCF and Cray initiated a hardware maintenance activity that removed no more than 20% of the compute blades at a time for repair. This strategy minimized the impact to the users while providing an adequate material supply to the rework process. The repair process was executed in three phases:

- at a Cray-subcontractor electronics manufacturing facility, where the primary connector mechanical assembly rework was completed;
- at Cray's manufacturing facility, where the rework was independently tested; and
- on-site at ORNL, where the reworked material was retested prior to reintroduction in to the production Titan partition.

This manufacturing and testing method resulted in very high confidence in the repair; that confidence has been demonstrated by subsequent node failure rates that are well within expected FIT rates for the system and with no connector repairs required for reworked blades. This maintenance activity was completed on December 17, 2013, at which time the full 18,688-node Titan partition was returned to service.

2.2.4 XK7 System Availability Summary

In addition to the planned downtime for this major system upgrade, the PCIe-3 repair and SXM mechanical assembly revisions adversely affected full system availability in 2013. Table 2.1 outlines the steps in the upgrade. Figure 2.1 summarizes the node availability during that time. Despite the significant scope of these activities, the OLCF exceeded its commitments to the INCITE, ALCC, and DD programs for Titan core-hours, delivering more than 2.6 billion compute hours among all programs in 2013, including staff programs, workshops, and other supporting interests. The OLCF staff and management have many decades of combined experience in fielding leadership systems. They can efficiently plan for known challenges and quickly and effectively respond to unexpected events. For these events, OLCF technical experts worked closely to resolve the hardware issues described, and OLCF managers communicated regularly with users and established policies and revised schedule priorities in order to maximize user access throughout the year.

Table 2.1. Time Line for the Cray XK7 Compute Blade and NVIDIA SXM Mechanical Assembly Revision

Date	Hardware Status
12/17/2013	18,688-node Cray Titan partition returned to production.
09/03/2013	NVIDIA SXM mechanical assembly rework begins. Available Titan partition guaranteed to meet or exceed 80% (16,384 nodes) of full system size.
05/31/2013	18,688-node Cray Titan released to production.
04/26/2013	Begin Titan acceptance test.
03/11/2013	8,972-node Cray Titan partition (96 cabinets) introduced to service. Full access to Kepler nodes enabled. PCIe-3 repairs on 1–4 cabinets.
02/02/2013	9,716-node Cray Titan partition (104 cabinets) remains in service. General availability restricted to AMD Opteron CPUs only. PCIe-3 repairs on 96 cabinets.
01/15/2013	18,688-node Cray Titan introduced to service. General availability restricted to AMD Opterons only.
10/08/2012	Cray JaguarPF removed from service.

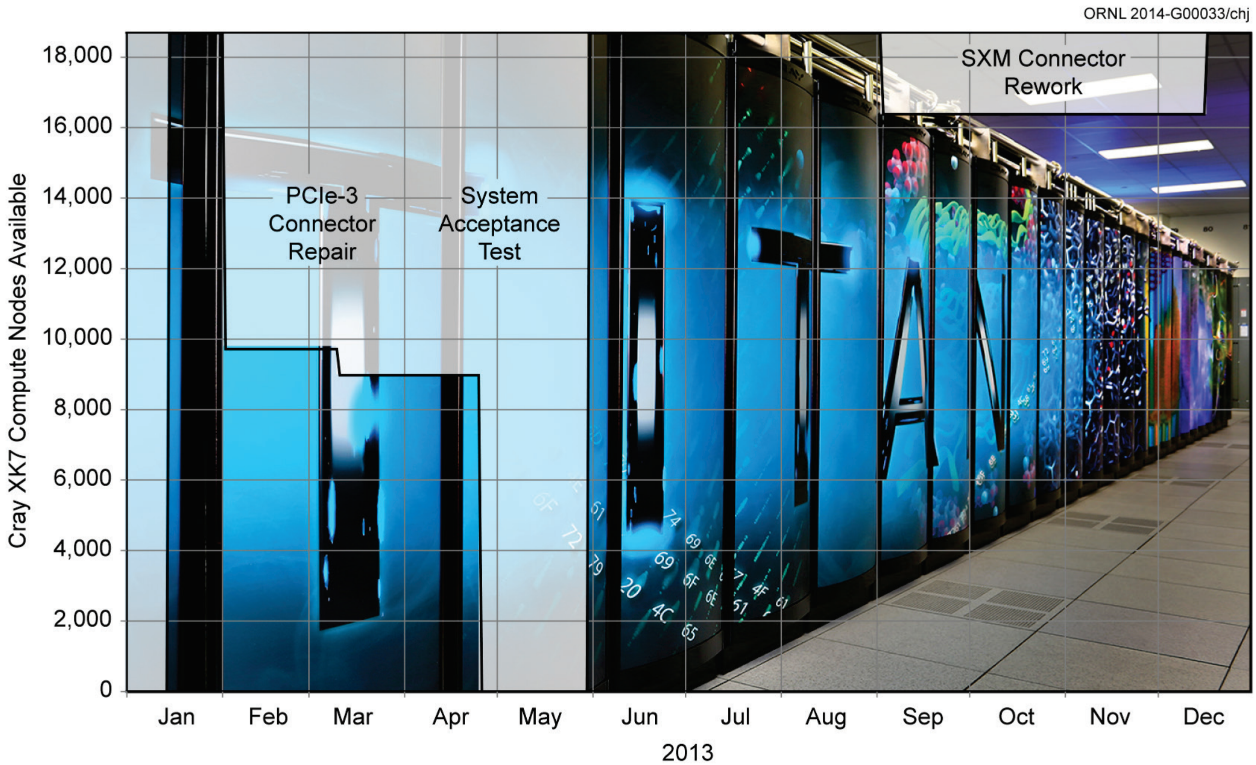

Figure 2.1. Cray XK7 Titan nodes available during 2013.

2.3 EOS—CRAY XC30

In September 2013, the OLCF installed and accepted a four-cabinet Cray XC30, called "Eos." The Eos system was released to production on October 3, 2013. All INCITE users were automatically granted access to the XC30. The system, with 744 Intel Xeon E5-2670 compute nodes and 47.6 terabybtes of memory, provides the OLCF user community with a substantive large-memory-per-node computing platform. The Eos nodes are connected by Cray's Aries interconnect in a network topology called "Dragonfly."

2.4 SPIDER II—ATLAS FILE SYSTEMS

In September, the OLCF released Spider II, its next-generation Lustre parallel file system, to production. Spider II is the architectural revision of the original Spider file system, which was installed in 2009 and which remain in production through instantiations of the /widow[*] file systems. Spider II contains two instantiations of the /atlas file system, with an aggregate capacity of more than 30 petabytes, and block-level performance of more than 1.3 terabyte/second. The /atlas file systems were mounted on all compute systems in the last quarter of 2013 and will be the default file systems for Titan beginning in the first quarter of 2014.

2.5 RHEA—DATA ANALYSIS AND VISUALIZATION CLUSTER

In the fourth quarter of 2013, the OLCF completed the installation and testing of a 196-node Linux cluster, called "Rhea". Rhea was released to production on January 8, 2014. This system, with 392 Intel E5-2650 processors and 12.5 terabytes of memory, provides a conduit for large-scale scientific discovery via data analysis and visualization of simulation data generated on Titan. Users with accounts on INCITE- or ALCC-supported projects were automatically provided access to Rhea; DD projects may also request access to Rhea. Rhea’s compute and I/O nodes are interconnected using fourteen-data-rate (FDR) InfiniBand; this same technology connects the cluster to the /atlas file systems.

2.6 EXPLORATORY VISUALIZATION ENVIRONMENT FOR RESEARCH IN SCIENCE AND TECHNOLOGY

Many of the scientific breakthroughs are enabled by the data analytics and visualization capabilities provided to the OLCF users. In 2013, the OLCF coordinated the complete redesign, deployment, and management of the Exploratory Visualization Environment for Research in Science and Technology (EVEREST) visualization laboratory as a state-of-the-art scientific discovery facility (Figure 2.2). Within EVEREST, there are three computing systems and two separate state-of-the-art visualization display walls. The primary display wall spans 30.5 feet × 8.5 feet and consists of 18 1920 × 1080 stereoscopic Barco projection displays arranged in a 6 × 3 configuration. The secondary display wall contains16 1920 × 1080 planar displays arranged in a 4 × 4 configuration, providing a standard 16:9 aspect ratio. With the new upgrade, scientists will be better able to visualize their data to make discoveries. The larger stereoscopic display provides 37 million pixels; the stereoscopic design creates depth in an image for a 3D effect. While the previous systems could show movies, high-resolution images, and some live applications, the new system’s 3D capabilities will allow for significantly greater detail in data visualization.

Figure 2.2. The EVEREST visualization laboratory.

2.7 OLCF COMPUTATIONAL RESOURCE SUMMARY

The OLCF provided the Titan and Eos computational resources in 2013 (see Table 2.2).

Table 2.2. OLCF Production Computer Systems, 2013

System	Access	Type	CPU	GPU	Computational Description			Interconnect
					Nodes	Node Configuration	Memory Configuration	
Titan	Full production	Cray XK7	2.2 GHz AMD Opteron 6274 (16-core)	732 MHz NVIDIA K20X (Kepler)	18,688	16-core SMP + 14 streaming multiprocessor (SM) GPU (hosted)	32 GB DDR3-1600 and 6 GB GDDR5 per node; 598,016 GB DDR3 and 112,128 GB GDDR5 aggregate	Gemini (Torus)
Eos	Full production	Cray XC30	2.6 GHz Intel E5-2670 (8-core)	None	744	2 × 8-core SMP	64 GB DDR3-1600 per node; 47,600 GB DDR3 aggregate	Aries (Dragonfly)

2.7.1 Null Time

The OA)business results are based on the length of time the computational resource has been in production. The results must consider the impact to the calculation of scheduled availability (SA), OA, mean time to interrupt (MTTI), mean time to failure (MTTF), and utilization due to "null time" for any system. A null time is a period of time within the reporting period that reduces the total potentially available time and is most frequently due to a system entering production after the beginning of the measurement period or to its leaving production before the end of the measurement period. In addition, null time may be attributed to systems when they must be removed from service for an extensive upgrade or similar activity. Null times are not considered in the calculation for SA, OA, MTTI, MTTF, and system utilization. The period of time devoted to the Titan acceptance is one example of null time associated with a system. The null times for OLCF systems in 2013 are noted in Table 2.3.

Table 2.3. OLCF HPC System Null Times, 2013

System	Null Time Start Date	Null Time End Date	Description
Computational Resources			
Titan	January 1, 2013 0:00	January 15, 2013 08:00	System not released to production
Titan	April 26, 2013 12:43	May 31, 2013 08:00	System acceptance testing
File Systems			
Atlas 0	January 1, 2013 0:00	October 3, 2013 8:00	Released to production during the year
Atlas 1	January 1, 2013 0:00	October 3, 2013 8:00	Released to production during the year
Widow 0[a]	March 5, 2013 8:00	December 31, 2013 23:59	System decommissioned

[a] Widow0 decommissioned on March 5, 2013.

2.7.2 OLCF HPC Resource Production Schedule

The OLCF production computational systems entered into production according to the schedule given in Table 2.4. This includes historical data associated with the Cray XT5, the very small overlap in December 2011 beginning with the introduction of the Cray XK6, and the series of Cray XK systems available in 2012 and 2013.

Table 2.4. OLCF HPC System Production Dates, 2008–Present

System	Type	Production Date[a]	Performance End Date[b]	Notes
Eos	Cray XC30	October 3, 2013	Null	744 Intel E5, 2,670 nodes.
Titan	Cray XK7	January 15, 2013 (initial) and May 31, 2013 (full)	Null	18,688 hybrid CPU-GPU nodes (AMD 6274/NVIDIA K20X).
JaguarPF	Cray XK6	September 18, 2012	October 7, 2012	Production at 240,000 cores until September 18, when partition size was reduced to 120,000 AMD Opteron cores. Additional Kepler installation. TitanDev access terminated.
JaguarPF	Cray XK6	February 13, 2012	September 12, 2012	Full production until September 12, when partition size was reduced to 240,000 AMD Opteron cores. Beginning of Kepler installation.
JaguarPF	Cray XK6	February 2, 2012	February 13, 2012	Stability Test. Restricted user access. 299,008 AMD Opteron 6274 cores. Includes 960-node Fermi-equipped partition.
JaguarPF	Cray XK6	January 5, 2012	February 1, 2012	Acceptance. No general access. 299,008 AMD Opteron cores.
JaguarPF	Cray XK6	December 12, 2011	January 4, 2012	142,848 AMD Opteron cores.
JaguarPF	Cray XT5	October 17, 2011	December 11, 2011	117,120 AMD Opteron cores.
JaguarPF	Cray XT5	October 10, 2011	October 16, 2011	162,240 AMD Opteron cores.
JaguarPF	Cray XT5	September 25, 2009	October 9, 2011	224,256 AMD Opteron cores.
JaguarPF	Cray XT5	August 19, 2008	July 28, 2009	151,000 AMD Opteron cores.

[a] The production date used for computing statistics is either the initial production date or the production date of the last substantive upgrade to the computational resource.
[b] The performance end date is the last calendar day that user jobs were allowed to execute on that partition.

For a period of 1 year following either system acceptance or a major system upgrade, the SA target for that HPC computational or storage system is at least 85% and the OA target is at least 80%.

2.7.3 Business Results Snapshot

Business results are provided for the OLCF computational resources, the High-Performance Storage System (HPSS) archive system, and the external Lustre file systems (see Tables 2.5, 2.6, and 2.7). Because an outage that may define the SA, OA, MTTI, or MTTF may occur outside the reporting period, the data reflected here artificially assume calculation boundaries of January 1, 2013 0:00 and January 1, 2014 0:00.

Table 2.5. OLCF Business Results Summary for HPC Systems

	Measurement	2012 Target	2012 Actual	2013 Target	2013 Actual
Cray XK7 (Titan)	Scheduled Availability	NIP	NIP	85.0%	98.70%
	Overall Availability	NIP	NIP	80.0%	93.82%
	MTTI (hours)	NIP	NIP	NAM	173.47
	MTTF (hours)	NIP	NIP	NAM	467.94
	Total Usage	NIP	NIP	NAM	89.93%
	Core-Hours Used	NIP	NIP	NAM	2,640,915,296
	Core-Hours Available	NIP	NIP	NAM	2,936,516,529
	Capability Usage				
	INCITE Projects	NIP	NIP	NAM	60.31%
	All Projects	NIP	NIP	30.0%	59.38%
Cray XE6/XK6 (JaguarPF)	Scheduled Availability	85.0%	98.11%	NIP	NIP
	Overall Availability	80.0%	91.45%	NIP	NIP
	MTTI (hours)	NAM	132.89	NIP	NIP
	MTTF (hours)	NAM	225.59	NIP	NIP
	Total Usage	NAM	84.39%	NIP	NIP
	Core-Hours Used	NAM	1,452,936,146	NIP	NIP
	Core-Hours Available	NAM	1,721,620,377	NIP	NIP
	Capability Usage				
	INCITE Projects	NAM	48.36%	NIP	NIP
	All Projects	30.0%	50.67%	NIP	NIP

MTTF = Mean time to failure.
MTTI = Mean time to interrupt.
NAM = Not a metric. No defined metric or target exists for this system. Data provided as reference only.
NIP = Not in production. This system was not available as a production resource.

Table 2.6. OLCF Business Results Summary for HPSS

	Measurement	2012 Target	2012 Actual	2013 Target	2013 Actual
HPSS	Scheduled Availability	95.0%	99.57%	95.0%	99.99%
	Overall Availability	90.0%	98.46%	90.0%	97.60%
	MTTI (hours)	NAM	228.6	NAM	450.00
	MTTF (hours)	NAM	588.85	NAM	2,919.78

MTTF = Mean time to failure.
MTTI = Mean time to interrupt.

Table 2.7. OLCF Business Results Summary for the External Lustre File Systems

	Measurement	2012 Target	2012 Actual	2013 Target	2013 Actual
Atlas 1	Scheduled Availability	NIP	NIP	85.0%	99.67%
	Overall Availability	NIP	NIP	80.0%	97.02%
	MTTI (hours)	NIP	NIP	NAM	232.97
	MTTF (hours)	NIP	NIP	NAM	430.82
Atlas 2	Scheduled Availability	NIP	NIP	85.0%	98.34%
	Overall Availability	NIP	NIP	80.0%	93.62%
	MTTI (hours)	NIP	NIP	NAM	224.79
	MTTF (hours)	NIP	NIP	NAM	425.35
Widow 0[a]	Scheduled Availability	NIP	NIP	85.0%	100.00%
	Overall Availability	NIP	NIP	80.0%	99.11%
	MTTI (hours)	NIP	NIP	NAM	761.15
	MTTF (hours)	NIP	NIP	NAM	1,536.00
Widow 1	Scheduled Availability	95.0%	99.88%	95.0%	99.87%
	Overall Availability	90.0%	98.25%	90.0%	97.35%
	MTTI (hours)	NAM	719.15	NAM	655.99
	MTTF (hours)	NAM	2,924.48	NAM	2,187.33
Widow 2	Scheduled Availability	95.0%	99.81%	95.0%	99.91%
	Overall Availability	90.0%	98.69%	90.0%	97.52%
	MTTI (hours)	NAM	722.42	NAM	657.14
	MTTF (hours)	NAM	2,191.89	NAM	2,188.11
Widow 3	Scheduled Availability	95.0%	99.89%	95.0%	99.87%
	Overall Availability	90.0%	98.95%	90.0%	98.09%
	MTTI (hours)	NAM	869.14	NAM	781.13
	MTTF (hours)	NAM	1,754.82	NAM	2,916.27

[a] Widow0 decommissioned on March 5, 2013.
MTTF = Mean time to failure.
MTTI = Mean time to interrupt.

2.8 RESOURCE AVAILABILITY

2.8.1 Scheduled Availability

2013 Operational Assessment Guidance

For HPC Facilities, scheduled availability (reference formula #1) is the percentage of time a designated level of resource is available to users, excluding scheduled downtime for maintenance and upgrades. To be considered a scheduled outage, the user community must be notified of the need for a maintenance event window no less than 24 hours in advance of the outage (emergency fixes). Users will be notified of regularly scheduled maintenance in advance, on a schedule that provides sufficient notification, and no less than 72 hours prior to the event, and preferably as much as seven calendar days prior. If that regularly scheduled maintenance is not needed, users will be informed of the cancellation of *that maintenance event in a timely manner. Any interruption of service that does not meet the minimum notification window is categorized as an unscheduled outage.*

A significant event that delays a return to scheduled production will be counted as an adjacent unscheduled outage. Typically, this would be for a return to service four or more hours later than the scheduled end time. The centers have not yet agreed on a specific definition for this improbable scenario.

$$SA = \left(\frac{time\ in\ period - time\ unavailable\ due\ to\ outages\ in\ period}{time\ in\ period - time\ unavailable\ due\ to\ scheduled\ outages\ in\ period} \right) * 100 \quad (1)$$

As shown in Table 2.8, the OLCF has exceeded the scheduled availability targets for the facility's computational resources for 2012 and 2013.

Table 2.8. OLCF Business Results Summary: Scheduled Availability

	System	2012 Target	2012 Actual	2013 Target	2013 Actual
Scheduled Availability	Cray XK7	NIP	NIP	85.0%	98.70%
	Cray XK6	85.0%	98.11%	NIP	NIP
	HPSS	95.0%	99.57%	95.0%	99.99%
	Atlas 1	NIP	NIP	85.0%	99.67%
	Atlas 2	NIP	NIP	85.0%	98.34%
	Widow 0[a]	NIP	NIP	85.0%	100.00%
	Widow 1	95.0%	99.88%	95.0%	99.87%
	Widow 2	95.0%	99.81%	95.0%	99.91%
	Widow 3	95.0%	99.89%	95.0%	99.87%

[a] Widow0 decommissioned on March 5, 2013.
NIP = Not in production. This system was not available as a production resource.

2.8.1.1 Assessing Impacts to Scheduled Availability

The operational posture for the Cray XK system contains a regularly scheduled weekly preventative maintenance (PM) period. PM is exercised only with the concurrence of the Cray hardware and software teams and with the HPC operations team. Typical PM includes software updates, application of field notices, and hardware maintenance to replace failed components. Without concurrence, the systems are allowed to continue operation.

2.8.2 Overall Availability

2013 Operational Assessment Guidance

Overall availability (reference formula #2) is the percentage of time a system is available to users. Outage time reflects both scheduled and unscheduled outages.

$$OA = \left(\frac{time\ in\ period - time\ unavailable\ due\ to\ outages\ in\ period}{time\ in\ period} \right) * 100 \quad (2)$$

As shown in Table 2.9, the OLCF has exceeded the overall availability targets for the facility's computational resources for 2012 and 2013.

Table 2.9. OLCF Business Results Summary: Overall Availability

	System	2012 Target	2012 Actual	2013 Target	2013 Actual
Overall Availability	Cray XK7	NIP	NIP	80.0%	93.82%
	Cray XK6	80.0%	91.45%	NIP	NIP
	HPSS	90.0%	98.46%	90.0%	97.60%
	Atlas 1	NIP	NIP	80.0%	97.02%
	Atlas 2	NIP	NIP	80.0%	93.62%
	Widow 0 [a]	NIP	NIP	80.0%	99.11%
	Widow 1	90.0%	98.25%	90.0%	97.35%
	Widow 2	90.0%	98.69%	90.0%	97.52%
	Widow 3	90.0%	98.95%	90.0%	98.09%

[a] Widow0 decommissioned on March 5, 2013.
NIP = Not in production. This system was not available as a production resource.

2.8.3 Mean Time to Interrupt

2013 Operational Assessment Guidance

Time, on average, to any outage on the system, whether unscheduled or scheduled. Also known as MTBI (Mean Time between Interrupt, reference formula #3).

$$MTTI = \left(\frac{time\ in\ period - (duration\ of\ scheduled\ outages + duration\ of\ unscheduled\ outages)}{number\ of\ scheduled\ outages + number\ of\ unscheduled\ outages + 1} \right) \quad (3)$$

where

time in period is start time–end time,

start time = end of last outage prior to reporting period,

end time = start of first outage after reporting period (if available) or start of the last outage in the reporting period.

The MTTI summary is shown in Table 2.10.

Table 2.10. OLCF Business Results Summary: Mean Time to Interrupt

	System	2012 Target	2012 Actual	2013 Target	2013 Actual
MTTI (hours)	Cray XK7	NIP	NIP	NAM	173.47
	Cray XK6	NAM	132.89	NIP	NIP
	HPSS	NAM	228.6	NAM	450.00
	Atlas 1	NIP	NIP	NAM	232.97
	Atlas 2	NIP	NIP	NAM	224.79
	Widow 0[a]	NIP	NIP	NAM	761.15
	Widow 1	NAM	719.15	NAM	655.99
	Widow 2	NAM	722.42	NAM	657.14
	Widow 3	NAM	869.14	NAM	781.13

[a] Widow0 decommissioned on March 05, 2013.
NAM = Not a metric. No defined metric or target exists for this system. Data provided as reference only.
NIP = Not in production. This system was not available as a production resource.

2.8.4 Mean Time to Failure

2013 Operational Assessment Guidance

Time, on average, to an unscheduled outage on the system (reference formula #4).

$$MTTF = \frac{time\ in\ period - (duration\ of\ unscheduled\ outages)}{number\ of\ unscheduled\ outages + 1} \quad (4)$$

where

time in period is start time–end time,

start time = end of last outage prior to reporting period,

end time = start of first outage after reporting period (if available) or start of the last outage in the reporting period.

The MTTF summary is shown in Table 2.11.

Table 2.11. OLCF Business Results Summary: Mean Time to Failure

	System	2012 Target	2012 Actual	2013 Target	2013 Actual
MTTF (hours)	Cray XK7	NIP	NIP	NAM	467.94
	Cray XK6	NAM	225.59	NIP	NIP
	HPSS	NAM	588.85	NAM	2,919.78
	Atlas 1	NIP	NIP	NAM	430.82
	Atlas 2	NIP	NIP	NAM	425.35
	Widow 0[a]	NIP	NIP	NAM	1,536.00
	Widow 1	NAM	2,924.48	NAM	2,187.33
	Widow 2	NAM	2,191.89	NAM	2,188.11
	Widow 3	NAM	1,754.82	NAM	2,916.27

[a] Widow0 decommissioned on March 5, 2013.
NAM = Not a metric. No defined metric nor target exists for this system. Data provided as reference only.
NIP = Not in production. This system was not available as a production resource.

2.9 RESOURCE UTILIZATION

2013 Operational Assessment Guidance

The Facility reports Total System Utilization for each HPC computational system as agreed upon with the Program Manager. This is reported as a number, not a metric.

Observation: The numbers that are reported for the Cray XK7 resource are Titan core-hours, where a single Titan node-hour comprises 16 AMD Opteron core-hours and 14 NVIDIA Kepler SM-hours. The OLCF refers to the combination of these traditional core-hours and SM-hours as "Titan core-hours", denoting that they are the product of a hybrid node architecture. Subsequent versions of this calculation may need to be revised to better reflect the specific systems at a particular Facility.

2.9.1 Resource Utilization Snapshot

For the Cray XK7 for the operational assessment period January 1–December 31, 2013, 2,640,915,296 Titan core-hours were utilized from an available 2,936,516,529 Titan core-hours. These numbers adjust for applicable null times, where hours consumed by the staff, the vendor, and other parties during a null time are removed from the calculation. This resulted in total system utilization for the Cray XK7 of 89.93%.

2.9.1.1 Resource Utilization Measurement Units

For 2013, system accounting was managed based on the characteristics of the node type that was used. For the period January 15, 2013–May 30, 2013, system accounting measured the consumption of 1 node-hour of computing time on Titan as 16 "Titan core-hours." This accounting decision was based on the fact that the Kepler GPU was not generally available and had not completed the OLCF acceptance test. Beginning May 31, 2013, system accounting was revised to measure the consumption of 1 Titan node-hour as 30 Titan core-hours (the combination of 16 Opteron core-hours and 14 NVIDIA Kepler SM-hours).

For 2013, allocations among all programs were provided to approved projects in terms of Titan core-hours. For 2014, INCITE allocations were provided to approved projects in terms of Titan core-hours; other 2014 programs will be allocated in terms of Titan node-hours, where there remains a direct correlation from the Titan node-hour to the Titan core-hour. This adjustment is consistent with the method required for allocating node-hours to applications. The job scheduler for the OLCF compute resources is Adaptive Computing's Moab, coupled to the Cray resource manager, Torque. Moab/Torque allocates resources at the granularity of a single node, not a core, regardless of the composition of that node. By migrating to an allocation program that consistently uses Titan node-hours, there is a coherent approach to system scheduling, accounting, and reporting.

Eos was prioritized as an additional resource for INCITE projects for the entire period of production in 2013. All INCITE projects were granted access to Eos. The charging factor for usage on Eos was 30 core-hours per node, the same as Titan. Through 2013, consumed hours by INCITE projects were counted toward INCITE usage, but were not subtracted from the project's INCITE allocation.

2.9.2 Total System Utilization

2013 Operational Assessment Guidance

The percent of time that the system's computational nodes run user jobs. No adjustment is made to exclude any user group, including staff and vendors (reference formula #5).

$$SU = \left(\frac{\text{core hours used in period}}{\text{core hours available in period}} \right) * 100 \qquad (5)$$

The Cray XK7 Titan system utilization is shown in Table 2.12. The measurement period is for 2013, irrespective of the prescribed allocation period of any single program. As an example, the INCITE allocation period follows a calendar year schedule. The ALCC program follows an allocation cycle that runs for 12 months beginning July 1 of each year. Utilization described here does not account for staff, vendor, or other time that may have been accrued during null times. System utilization for 2013 is 89.93%.

Table 2.12. 2013 Cray XK7 Titan Utilization

Time Period	CPU Hours Consumed	CPU Hours Available	Percent of Allocation Consumed
January	82,395,654	116,010,121	71.02%
February	98,392,378	105,214,518	93.52%
March	93,831,227	101,489,977	92.45%
April	89,081,222	115,088,862	77.40%
May	6,707,157	8,970,240	74.77%
June	345,839,232	391,364,096	88.37%
July	353,135,832	401,474,304	87.96%
August	355,347,891	401,483,648	88.51%
September	294,229,130	326,470,440	90.12%
October	273,686,028	292,228,640	93.65%
November	290,101,839	304,884,800	95.15%
December	358,167,707	371,836,883	96.32%
Total	2,640,915,296	2,936,516,529	89.93%

The OLCF tracks the consumption of Titan core-hours by job. This can then be extended to track with high fidelity the consumption of Titan core-hours by program, project, user, and system. Figure 2.3 summarizes the Cray XK7 utilization by month and by program for all of 2013. No adjustment is made to exclude any user group, including staff and vendors. Figure 2.3 additionally describes preproduction activity through May 30, 2013.

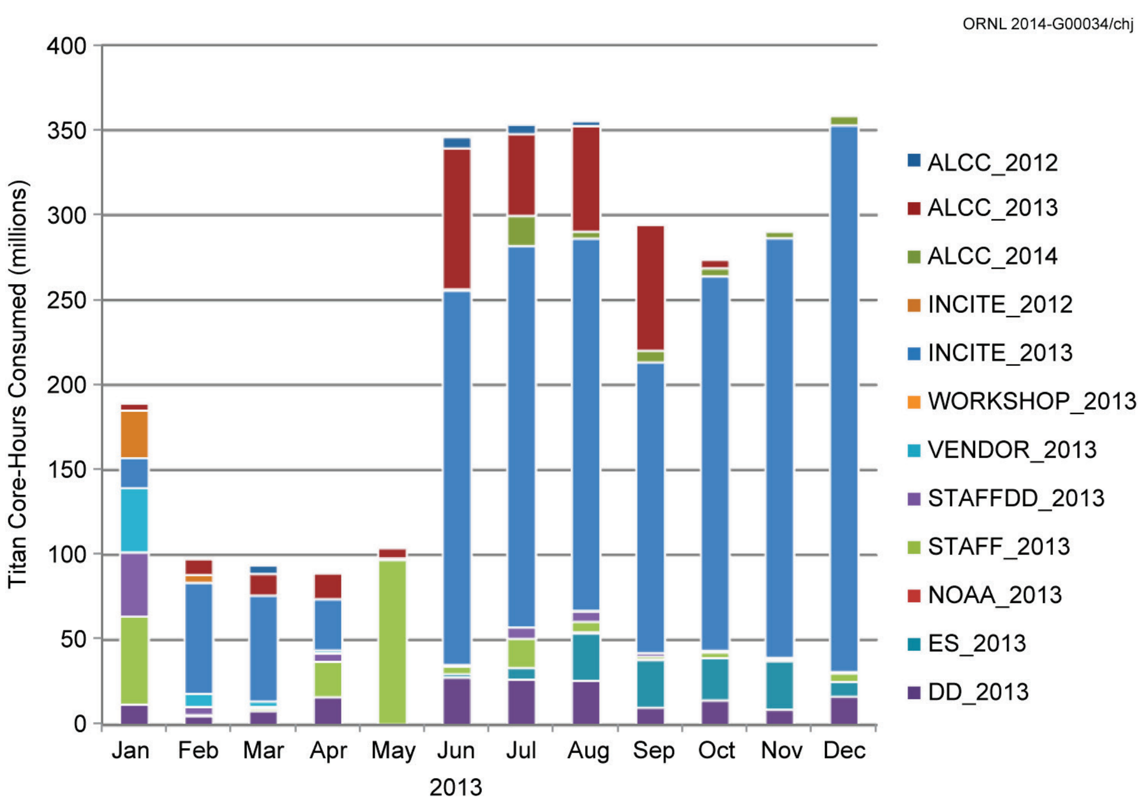

Figure 2.3. 2013 XK7 Resource Utilization—Titan Core-Hours by Program.

2.9.2.1 Assessing Total System Utilization

The monthly utilization of available Titan core-hours is shown in Figure 2.4. The reduced availability in the first half of 2013 due to the PCIe-3 connector repair, the SXM connector repair (September–December), and the null time with the acceptance test in April and May did not diminish demand for the system; utilization of the available core-hours remained very high.

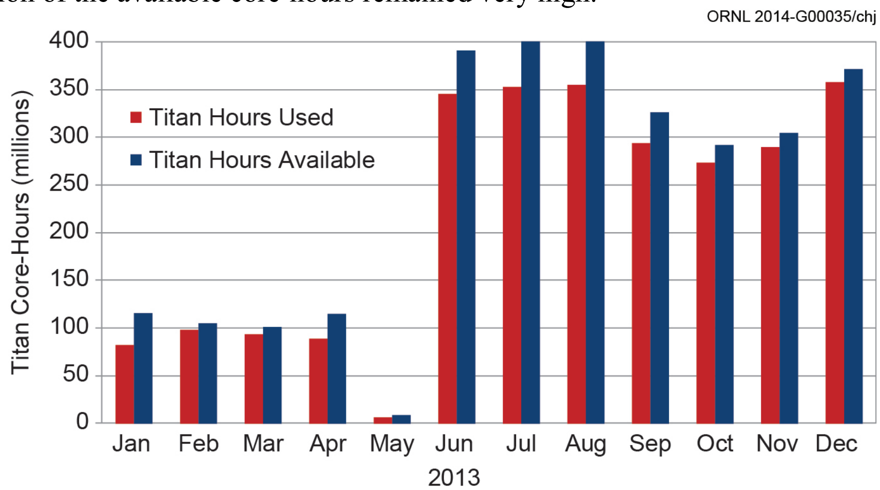

Figure 2.4. High System Utilization on Titan in 2013.

2.9.2.2 Performance of the Allocated Programs

All allocation programs, including INCITE, ALCC, and DD, are aggressively monitored to ensure that projects within these allocation groups maintain adequate consumption rates. The 2013 INCITE allocation program was the largest program in 2013 with a commitment for 1.84B Titan core-hours. The consumption of these allocation programs is shown in Table 2.13.

Table 2.13. 2013 Allocated Program Performance on the OLCF Resources

Program	Allocation	Hours Consumed	Percent of Total
INCITE	1,840,000,000	1,944,595,949	73.97%
Titan		1,889,570,654	
Eos		55,025,295	
ALCC	Not applicable	384,837,566	14.64%
ALCC_2012		21,506,184	
ALCC_2013		320,068,950	
ALCC_2014		43,262,432	
DD	Not applicable	171,404,416	6.52%
ES	Not applicable	126,543,799	4.81%
NOAA	Not applicable	1,518,621	0.06%
Total		2,628,900,351	100.00%

INCITE = Innovative and Novel Computational Impact on Theory and Experiment; period of performance reported: January 1, 2013-February 5, 2014
ALCC = Advanced Scientific Computing Research Leadership Computing Challenge; period of performance is July 1 through June 30 of the following year (e.g., ALCC_2014 is July 1, 2013 through June 30, 2014). Numbers reported here are for any ALCC project active during 2013.
DD = Director's Discretionary; period of performance reported: January 1, 2013-December 31, 2013
ES = Early Science; period of performance reported: January 1, 2013-December 31, 2013
NOAA = National Oceanic and Atmospheric Administration; period of performance reported: January 1, 2013-December 31, 2013

Note that non-renewed INCITE projects from 2012 were allowed by OLCF policy to continue running at low priority early in the first quarter of 2013 so that those projects could complete while new 2013 INCITE projects ramped up. Not only is this a user-friendly policy for non-renewed projects that have not quite exhausted their allocation, it serves to increase utilization while new projects establish a more predictable consumption routine. This policy remained in effect for this reporting period; 2013 non-renewing INCITE projects that had not exhausted their allocation were allowed to continue to submit jobs in accordance with existing queue policy to the workload managers on both Titan and Eos through January 31, 2014. ALCC projects from the 2012 and 2013 allocation periods were granted extensions as well. These policies allowed those programs to more fully use their allocations despite challenges associated with these nascent programs moving to the very large system or the upgrade schedule that reduced access to the system.

2.10 CAPABILITY UTILIZATION

2013 Operational Assessment Guidance—Capability Utilization

The Facility shall describe the agreed definition of capability, the agreed metric, and the operational measures that are taken to support the metric.

Capability usage defines the minimum number of nodes allocated to a particular job on the OLCF computing resources. To be a capability job, any single job must use at least 20% of the available nodes of the largest system (Titan).

The metric for capability utilization describes the aggregate number of node-hours delivered by capability jobs. In 2011, the metric stipulated that no less than 35% of the delivered node-hours on the Cray XT5 would reflect capability jobs. For the first year of Cray XK production (2012), the metric stipulated that no less than 30% of the delivered node-hours reflect capability jobs. As Titan reached production as a new system in 2013, the applicable metric is again 30%. This is proposed to increase to 35% in subsequent production years, beginning May 31, 2014.

The OLCF Resource Utilization Council uses queue policy on the Cray systems to support delivery of this metric target, providing queues specifically for capability jobs with 24-hour wall-clock times and increased priority.

The OLCF Capability Utilization Definition is summarized in Table 2.14.

Table 2.14. OLCF Capability Utilization Definition

System	Year 1		Subsequent Years	
	Definition for Capability	Capability Metric	Definition for Capability	Capability Metric
Cray XK7 Titan	20% of available nodes on the largest system	30% of delivered hours	20% of available nodes on the largest system	35% of delivered hours

The OLCF continues to exceed expectations for capability usage of its HPC resources (Table 2.15). Keys to the growth of capability usage include the liaison role provided by the SciComp Group members, who work hand-in-hand with users to port, tune, and scale code, and OLCF support of the application readiness efforts (CAAR), where staff actively engage with code developers to promote application portability, suitability to hybrid node systems, and performance. The OLCF aggressively prioritizes capability jobs in the scheduling system.

Table 2.15. OLCF Capability Usage on the Cray XK Systems

	Leadership Usage	CY2012 Target	CY2012 Actual	CY2013 Target	CY2013 Actual
Cray XK7	INCITE	NIP	NIP	NAM	60.31%
	Total	NIP	NIP	30%	59.38%
Cray XK6	INCITE	NAM	48.36%	NIP	NIP
	Total	30%	50.67%	NIP	NIP

NAM = Not a metric. No defined metric nor target exists for this system. Data provided as reference only.
NIP = Not in production. This system was not available as a production resource.

The average consumption of hours by capability jobs was well above the 2013 target of 30% at 59.38%. This consumption varies during the year, affected by factors including system availability and the progress by the various projects within their research. The distribution of the consumption of hours by capability jobs, by month, is shown in Figure 2.5.

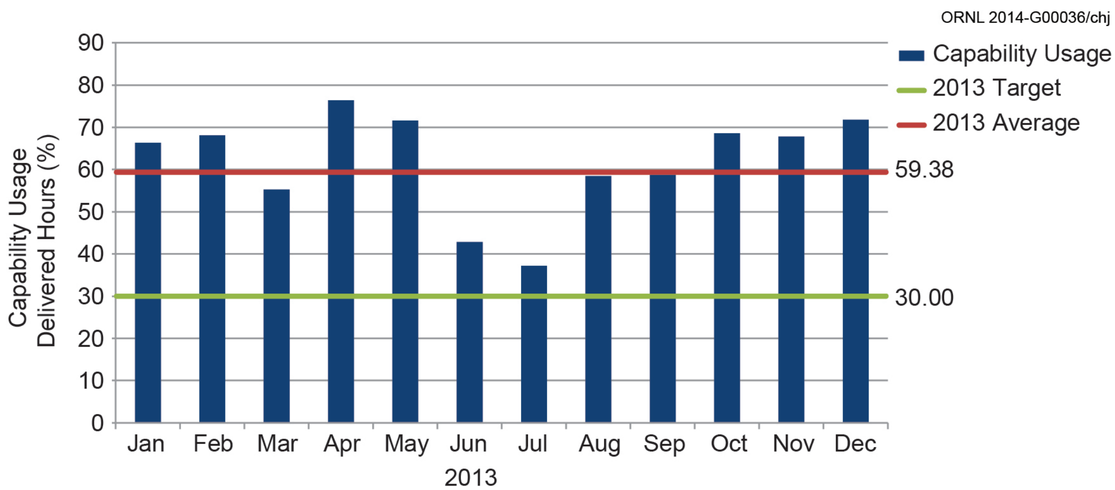

Figure 2.5. Effective Scheduling Policy Enables Capability Usage.

Capability jobs are not restricted to the INCITE program. There are capability jobs across the Early Science, ALCC, and DD programs as well. The contribution to capability utilization by program is shown in Figure 2.6.

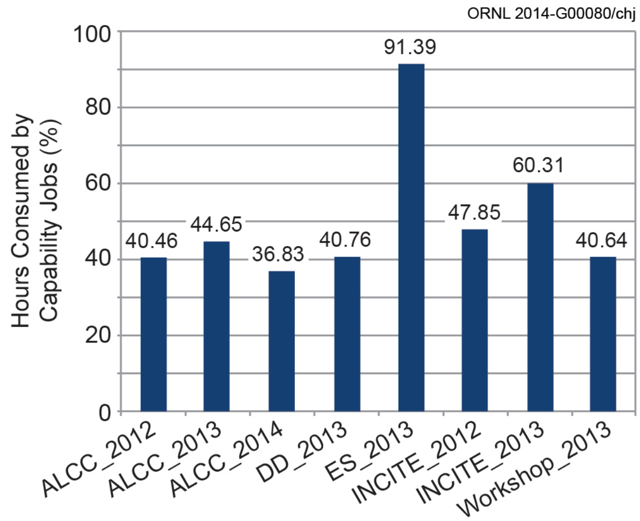

Figure 2.6. Capability Usage by Project Type.

2.11 GPU USAGE

The upgrade of the existing Cray XT5 to the XK series system provided a key new capability to users, allowing them to exploit a new hybrid compute node that contains both a CPU and an NVIDIA accelerator. On any hybrid node, the GPU is an option for the user. There is no explicit requirement to use it and a hybrid node can be used in the exact same manner as an Opteron-only node.

2.11.1 Measuring GPU Usage with Automatic Tracking Library Database

For 2013, the OLCF continued to rely upon the Automatic Library Tracking Database (ALTD), which was first used in 2012, as a method for determining whether an application was employing the GPU. ALTD actively monitors the compilation phase of individual applications and at link time creates a unique record for that application that contains a list of each of the libraries that were linked against that particular binary. When this application is executed via aprun, a new ALTD record is written to the database that contains the name of the executable, the batch job ID, and other supporting information. To determine whether a specific executable takes advantage of the GPU, we examine whether an executed job, for which we have all of the per-job scheduling information, was linked against an accelerator-specific library. For 2013, this includes any library that matches the following identifiers:

libacc*, libOpenCL*, libmagma*, libhmpp*, libcuda*, libcupti*, libcula*, libcublas*

Jobs whose executables are linked against one of the above are deemed to have used the accelerator. From this information, per-job utilization can be derived and aggregated into reports that describe system utilization across GPU-enabled and CPU-only qualifiers.

While this method does provide a mechanism for examining CPU-only and GPU-accelerated contributions to system utilization, there are some limitations.

- While ALTD is enabled by default, it must be disabled in certain instances where there is a software stack conflict. Then, job compilations executed without inclusion in ALTD will not produce corresponding records, and subsequent execution of that binary will return a null lookup result.

- Debugging sessions cannot generally tolerate the wrapped aprun, so these sessions will contribute to an unknown result.

- Jobs that are executed outside the job scheduler, such as in a dedicated mode, will not generate job records and cannot be correlated. They will contribute to an unknown result.

The results of the implementation of this method are shown in Figure 2.7, where the information is aggregated on a monthly schedule and the results are expressed as the percentage of Titan core-hours that were delivered by applications using the GPUs.

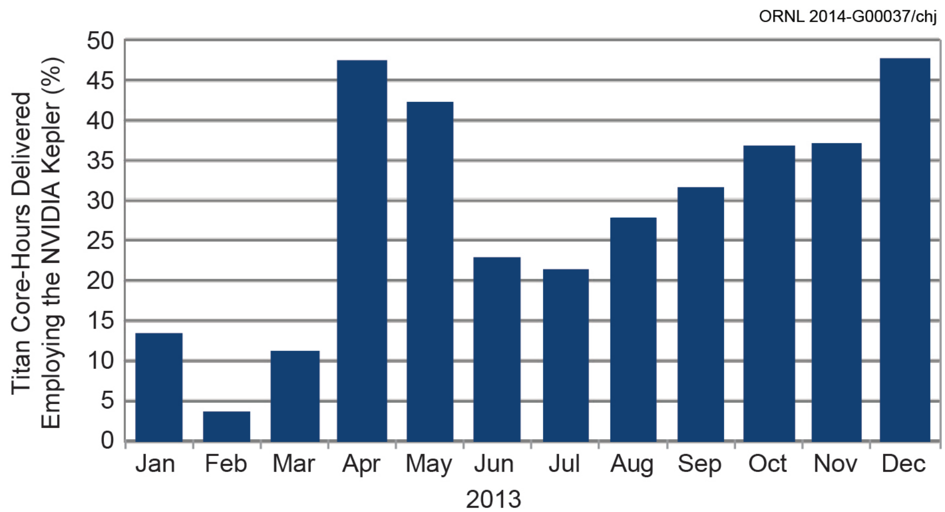

Figure 2.7. Tracking GPU Usage on Titan.

2.11.2 Accessing Limited NVIDIA Management Library Data through Cray's Resource Utilization Reporting

The method using ALTD makes a binary assumption: the application uses the GPU, or it does not. To understand the degree to which the GPU is being utilized requires more information directly from the GPU itself.

To meet this requirement, the OLCF actively engaged with both NVIDIA and Cray in 2012 and 2013 to define what information was needed from a system accounting perspective. The results of those discussions drove revisions to the NVIDIA device driver and to the accompanying API and library, and changes to the Cray Resource Utilization software so that additional information about the GPU usage, on a per-job basis, will be available at the conclusion of each job. This development effort is nearing conclusion, with the release of the updated NVIDIA driver and with changes to the available Cray software stack, beginning with CLE4.2 UP02. This software version was installed on the single-cabinet Cray XK7 in Q4CY14, and per-job GPU accounting information is being collected from that system. CLE 4.2 UP02 will be installed on Titan in Q1CY14; the job-accounting database has already been revised to accommodate the GPU accounting data.

2.12 SAFETY

The provision of a safe working environment and the demonstrated safety-conscious attitude of all subcontractors and employees remain important considerations. In the face of the very high volume of work required by the Cray XK7 maintenance activity, the ability to foster and promote a safe work environment remains paramount. See Section 6, "Site Office Safety Metrics," for a description of how the facility incorporates DOE site office safety recommendations into its operations.

There was a single reportable incident, on October 28, 2013, categorized as an SC3, 2E(2)—"Unexpected Discovery of an Uncontrolled Electrical Hazardous Energy Source." This criterion does not include discoveries made by zero-energy checks and other precautionary investigations made before work is authorized to begin.

Category 3 SC-ORO--ORNL-X10EAST-2013-0010, “Discovery of an Unexpected Energy Source after Replacement of a Circuit Relay” was submitted to the DOE Occurrence Reporting and Processing System on October 30, 2013. The Assessment and Commitment Tracking System (ACTS) number for this issue is 0.29220. Reference is https://orps.hss.doe.gov/orps/reports/displayReport.asp?idx=132364.

Description of Occurrence: On October 28, 2013, a qualified electrical worker replaced a 120 volt circuit relay after performing Lock-Out/Tag-Out (LOTO) at the breakers indicated on drawings. A zero energy check was performed at the relay’s electrical contacts indicating zero voltage. The LOTO and zero energy check were performed in accordance with ORNL procedures including the appropriate personal protective equipment. After establishing electrical safe working conditions, the worker removed his voltage rated gloves. The worker installed the new relay and performed post-maintenance testing which indicated abnormal results. Investigation identified the presence of an unexpected energized source. Line Management was immediately notified. ORNL line management notified the Laboratory Shift Superintendent and the event was categorized as an SC3, 2E(2) “Unexpected Discovery of an Uncontrolled Electrical Hazardous Energy Source.” There was no injury with hazardous energy or other environmental, health and safety impacts as a result of the event.

3. STRATEGIC RESULTS

CHARGE QUESTION 3: Is the facility enabling scientific achievements consistent with the Department of Energy strategic goals?

OLCF RESPONSE: Yes. The center continues to enable high-impact science results through access to the leadership-class systems and support resources. The allocation mechanisms are robust and effective.

The projects and user programs operating within the OLCF will advance the DOE's mission to ensure America's security and prosperity by addressing its energy, environmental, and nuclear challenges through transformative science and technology solutions. The modest number of accomplishments that are described in this section serve to communicate how the OLCF is advancing two of DOE's four strategic goals, and associated targeted outcomes, as stated in the *U.S. Department of Energy Strategic Plan* (May 2011):

- Goal 1: Catalyze the timely, material, and efficient transformation of the nation's energy system and secure U.S. leadership in clean energy technologies.
- Goal 2: Maintain a vibrant U.S. effort in science and engineering as a cornerstone of our economic prosperity with clear leadership in strategic areas.

3.1 SCIENCE OUTPUT

2013 Operational Assessment Guidance

The facility tracks and reports the number of refereed publications written annually based on using the facility's resources. This number may include publications in press or accepted but not submitted or in preparation. This is a reported number, not a metric. In addition, the facility may report other publications where appropriate.

3.1.1 OLCF Publications Report

Two hundred sixty-two refereed publications resulting from the use of OLCF resources were published in 2013 as identified in the search and compilation completed on February 4, 2014.[§] Only publications appearing in print were counted. In the 2012 OLCF OAR, 321 publications were reported; however, unlike the 2013 total, the total for 2012 included publications that at year's end were in press or accepted for publication.

The OLCF follows the recommendation from the 2007 report of the ASCR Advisory Committee Petascale Metrics Panel to report and track user products, including, for example, publications, project

[§] In this document, "year" refers to the calendar year unless it carries the prefix "FY," indicating the fiscal year.

milestones (requested quarterly; also examined in the INCITE renewal process), and code improvement. The methodology used for publication identification in 2013 is summarized in Section 3.1.2.

Sponsor guidance allows the reporting of "publications in press or accepted." In previous years, the OLCF has followed this guidance inclusively, reporting manuscripts accepted for publication but not appearing in print. However, papers that are accepted for publication but not yet generally available to the public will not be automatically discovered by search tools. Moreover, reporting papers that are not yet in print in 2013 but that will be available in 2014 may complicate the tracking of publications year over year. The OLCF considers this undesirable, and, therefore, the interpretation of the guidance has been modified so that only publications appearing in print in the year under review (e.g., 2013) are eligible for tabulation in the current report. The number of publications reported within previous OARs will be reevaluated in the light of this new interpretation with updates being communicated within regular reporting opportunities.

3.1.2 Methodology for OLCF Publication Discovery and Reporting

The OLCF requires an effective and efficient process to discover, curate, and report publications that have utilized the facility's computational resources, and no single collection method, if implemented in isolation, is sufficient to the task (e.g., user-reported lists of publications are error prone, and automated searches can be no more complete than the completeness of the database searched). However, tools developed at ORNL for automated discovery of documents from distributed databases have advanced the ability of the OLCF to reliably perform this task. Therefore, active database searches are integrated with the collection of user-reported publications to create a collection of publications to meet sponsor guidance for reporting.

Four data sets utilizing different methodologies for collection were generated and used to obtain the results for the 2013 *OLCF Operational Assessment Report*:

- An automated discovery process, using a software tool named COBRA, applied to searching the Institute for Scientific Information (ISI) Web of Science publication database for papers published in 2013 by OLCF users;

- self reports of papers published in 2013 by OLCF users and ORNL staff supporting the OLCF program;

- access to ORNL's Publication Tracking System (PTS) for papers published in 2013 from OLCF users and ORNL staff; and

- A manual search for papers published in 2013 by OLCF users within the ISI Web of Science database was performed by the ORNL Research Library, and limited to the publications of the American Chemical Society (ACS), only (this is required because ACS prohibits the "crawling" of its online journals.)

- Table 3.1 contains statistics from the results of these four approaches. Of the more than fifty thousand publications retrieved by the four search methods, about one percent deemed potential OLCF publications. After all of the data sets were compiled into one, duplicates were removed, yielding 262 unique, confirmed publications that resulted from utilization of OLCF resources.

Table 3.1. Summary of Statistics from OLCF Publication-Discovery Methodology

Process	Publications Retrieved	Potential OLCF Publications	Unique, Confirmed OLCF Publications
COBRA	49,753	159	**120**
Self Reports	240	238	**104**
PTS	106	75	**24**
Library	1790	29	**14**
Total	**51,889**	**501**	**262**

3.2 SCIENTIFIC ACCOMPLISHMENTS

The OLCF advances DOE's science and engineering enterprise through robust partnerships with its users. The following subsections provide brief summaries of selected scientific and engineering accomplishments, as well as resources for obtaining more information. While they cannot capture the full scope and scale of achievements enabled at the OLCF in 2013, these accomplishments advance the state of the art in science and engineering research and development (R&D) and are advancing DOE's science programs toward their targeted outcomes and mission goals. As an additional indication of OLCF achievements, OLCF users published many breakthrough publications in high-impact journals in 2013, including one in *Nature*, two in *Nature Communications*, two in *Nature Scientific Reports,* three in *Geophysical Review Letters,* and 12 in *Physical Review Letters.*

3.2.1 The Bleeding "Edge" of Fusion: C. S. Chang, Princeton Plasma Physics Laboratory, INCITE

Objective: To use very large-scale simulations to better understand the formation of turbulent flow in the components of a fusion-reactor plasma—a roiling collection of electrons, protons, and neutral particles formed by the high temperatures generated from nuclear fusion—and the interplay between these flows and the fusion reactor vessel.

Impact: By using the particle-based methods developed by the researchers, coupled with the computing power of Titan, the project was able to identify the physical causes of and the subsequent effects of important interactions between the edge of the fusion plasma and the reactor walls.

The performance of a tokamak fusion reactor is greatly affected, and significantly controlled, by events at the plasma's edge. When the plasma comes into contact with the vessel wall, it loses mass and energy and introduces neutral particles back into the plasma. As a result, equilibrium physics generally do not apply at the edge, and simulating the environment is not possible using conventional computational fluid-dynamics approaches. In addition, formations of strong density fluctuations (or clumps) flow together and move large amounts of edge plasma around—in a process known as "blobby" turbulence—greatly affecting edge and core plasma performance.

Accomplishments: Edge plasma transport processes will determine the fusion efficiency of the ITER fusion science reactor, currently under construction in France. With Titan and C.S. Chang's XCG1 simulation code, a product of the Center for Edge Plasma Physics SciDAC project, first-principles simulations of edge physics are now possible and can be used to increase understanding of ITER, and also today's fusion reactors, such as D-IIID and JET. The objective of the work described here is to perform the first edge-plasma simulation of the D-IIID reactor at the General Atomics facility in San Deigo, California, using first-principles gyrokinetic modeling techniques.

The researchers used Titan to study the multiscale interaction among gyrokinetic ions, drift-kinetic electrons, and neutral particles in a realistic DIII-D tokamak geometry, including the magnetic X-point and the grounded wall. Importantly, they identified nonlinear coherent turbulence structures ("blobs") in the plasma edge of the DIII-D reactor, the momentum source of these structures, the inward transport process and were able to predict the divertor heat load distribution, all for the first time in first-principles calculations. Physical mechanisms driving these phenomena were identified. Success with DIII-D plasma simulations has given confidence for JET- and ITER-scale campaigns within the new INCITE program year. In a conversation with the NCCS director of science, the principal investigator, C.S. Chang, noted, "The achievements . . . should be credited to OLCF only. We could not run such an extreme scale job at NERSC."

The plasma density fluctuations relative to a constant density are illustrated in Figure 3.1. These density fluctuations are driven by turbulence originating at the plasma edge spreading inward while interacting with the background density that is being driven by the central plasma heat source. Eventually, the whole volume becomes turbulent, with the spatial turbulence amplitude distribution being just enough to produce the outward heat transport to expel the centrally deposited heat to the edge. The edge turbulence source is continuously fed by the heat flux from the core. This is how the plasma profile, the heat source and the turbulence self-organize.

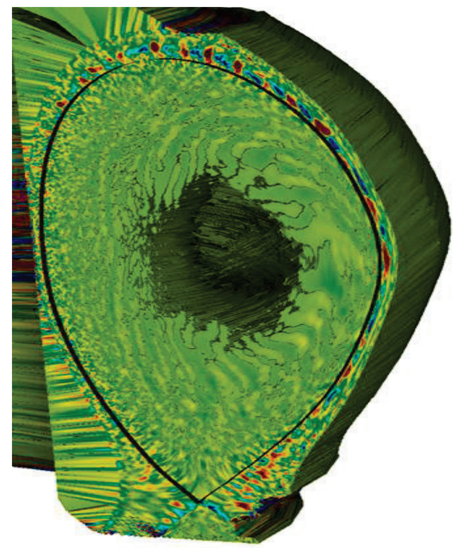

Figure 3.1. Plasma density fluctuations relative to a constant density. Visualization by David Pugmire (ORNL).

OLCF Contributions: Working with Ed D'Azevedo (supported by the OLCF and by the SciDAC project Center for Edge Physics Simulation), this team optimized its XGC1 code for Titan's GPUs using the maximum number of nodes, boosting performance fourfold over the previous CPU-only code. Relative to performance on the CPU-only Cray XK6 architecture, XGC1 performs 2.5× faster on Titan. David Pugmire wrote a data format reader and analyzed and visualized the simulation data. I/O and data workflow were managed through the ADIOS framework. This achievement required runs not possible at any other HPC center: 46% of work required 11,250 or more of Titan's XK7 nodes; 93% of the 102 million hours used was consumed at the capability scale (i.e., greater than 3,750 nodes).

Invited Talks:

C.S. Chang, "Extreme Scale Gyrokinetic Simulation of Multiscale Edge Blobs and its Implication to Edge-Core Interaction," 23rd International Toki Conference, Toki-City, Japan, November 18–21, 2013.

C.S. Chang, "What Can Magneto-Fluid Codes Learn from Full-f Gyrokinetics?" Workshop on Modeling Kinetic Aspects of Global MHD Modes Lorentz Center, Leiden, Netherlands, Dec. 2–6, 2013.

Online Story:

OLCF Staff Writer, "The Bleeding Edge of Fusion Research" *OLCF News* (February 14, 2014).

3.2.2 Physics-Based Probabilistic Seismic Hazard Analysis: Thomas Jordan, University of Southern California, INCITE

Objectives: To perform a probabilistic hazard analysis in order to calculate physics-based earthquake ground motions up to 10 Hz for use by engineers designing and building earthquake-resistant infrastructure. Ultimately, the aim is to calculate a California-wide CyberShake seismic hazard model using the Uniform California Earthquake Rupture Forecast.

Impact: The creation of the California-wide hazard model with high-frequency modeling was only possible through the team's access to Titan, where performance improvements to critical seismic simulation codes, coupled with improved workflow management, were implemented. These performance improvements also allow the CyberShake team to model and characterize ground motions at frequencies higher than 1 Hz, where the ground itself can attenuate, or reduce, the amplitude of the motion. Though Los Angeles has not experienced a major quake in its time as a city, the proximity of the San Andreas Fault provides a clear danger of such a quake occurring in the future. Data for the region is available from smaller quakes, but such information does not give emergency officials and structural engineers the information they need to prepare for a quake of magnitude 7.5 or greater.

Accomplishments: The calculations on Titan are milestones in the use of the GPU-accelerated code AWP-ODC to work toward CyberShake hazard maps above 1 Hz. This project integrated new physics into its ground motion modeling software, optimized its software to improve performance at large scale, and then used the software to run high-frequency deterministic simulations of scenario earthquakes at frequencies of interest to engineering studies.

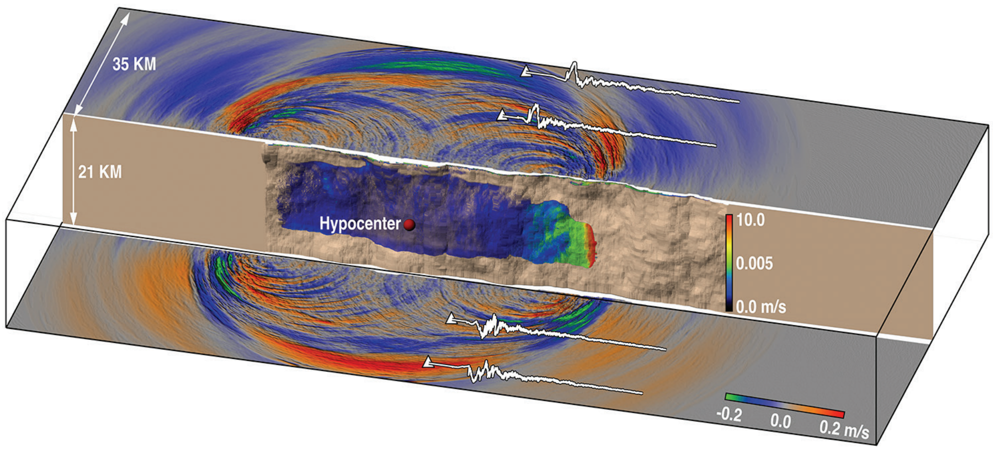

Figure 3.2. Snapshots of 10-Hz rupture propagation (slip rate) and surface wavefield (strike-parallel component) for a crustal model with a statistical model of small-scale heterogeneities. Visualization courtesy of the Southern California Earthquake Center.

New physics was integrated into the AWP-ODC-GPU software to model the effects of frequency-dependent attenuation of ground motions with distance in order to produce more accurate ground motion results at higher frequency (Figure 3.2). Frequency-dependent attenuation was considered negligible for simulations up to 1 Hz but important to higher-frequency simulations. The code was then optimized for use with Titan's GPUs, reaching a sustained performance of 2.3 petaflops on 16,384 Titan nodes and enabling the first-ever 10 Hz deterministic wave propagation simulations. Performance improved by a factor of 3.7 on the Titan XK7 compared to the dual-CPU Blue Waters Cray XE6 system. Speedup by a factor of 3.7 on accelerated Titan XK7, compared to the heavily optimized CPU code running on XE6, was achieved in key strain tensor calculations critical to probabilistic seismic hazard analysis. Then, the project ran deterministic, three-dimensional ground motion simulations representative of a 7.2-magnitude quake in the Los Angeles region, with maximum simulated frequencies up to 10 Hz. Results are now being compared to ground motion recordings from well-observed earthquakes. Once validated, the improved high-performance AWP-ODC-GPU code will be used for reciprocity-based probabilistic CyberShake seismic hazard model calculations.

OLCF Contributions: This achievement required runs not possible at any other HPC center: 93% of the 42 million hours used was consumed at a scale greater than 60% of Titan, specifically,16,640 nodes, and 95% of this project's computing utilized Titan's GPU accelerators.

Related Publications:

Y. Cui, E. Poyraz, K. Olsen, J. Zhou, K. Withers, S. Callaghan, J. Larkin, C. Guest, D. Choi, A. Chourasia, Z. Shi, S. Day, P. Maechling, and T. Jordan, "Physics-based Seismic Hazard Analysis on Petascale Heterogeneous Supercomputers," Technical Paper, SC13, Denver, Nov 17–22, 2013.

J. Zhou, Y. Cui, E. Poyraz, D, Choi, Dong, C. Guest, "Multi-GPU Implementation of a 3D Finite Difference Time Domain Earthquake Code on Heterogeneous Supercomputers," 13th Annual International Conference on Computational Science (ICCS). 13th INTERNATIONAL CONFERENCE ON COMPUTATIONAL SCIENCE Book Series: Procedia Computer Science; Vol. 18, 1255–1264 (2013).

Online Stories:

OLCF Staff Writer, "Titan Simulates Earthquake Physics Necessary for Safer Building Design," *OLCF News* (December 16, 2013).

R. Brueckner, "Earthquake Simulations on Titan Make for Safer Buildings," *Inside HPC* (December 19, 2013).

J. Zverina, "UC San Diego Team Achieves Petaflop-Level Earthquake Simulations on GPU-Powered Supercomputers," *News Center*, The University of California, San Diego (April 2013).

C. Cole, "CUDA Spotlight: Yifeng Cui," *NVIDIA CUDA Spotlight* (July 3, 2013).

C. Cole, "CUDA Spotlight: GPU-Accelerated Earthquake Simulations," *NVIDIA CUDA Spotlight* (July 7, 2013).

Award Recognitions:

International Data Corporation, "IDC Announces New Winners of HPC Innovation Excellence Awards," *IDC Press Release* (June 18, 2013).

International Data Corporation, "IDC Announces Winners of Sixth HPC Innovation Excellence Awards," *IDC Press Release* (November 19, 2013).

3.2.3 Non-Icing Surfaces for Cold Climate Wind Turbines: M. Yamada, GE Global Research, ALCC

Objective: Despite the growing importance of wind energy— By 2020, it will provide as much as 12 percent of the planet's electricity supply— this source of clean energy is underused in cold climates, where ice on turbines reduces their efficiency and at times forces them to be shut down. Reducing wind turbine downtime in cold climates requires an understanding of the microscopic mechanism of water droplets freezing on surfaces and a determination of the efficacy of nonicing surfaces at different temperatures.

Impact: The project delivered the most comprehensive molecular dynamics calculation of water freezing on a surface ever performed. These results can be compared immediately to experimental results, forming a powerful synergy between computation and experiment that can be applied to future designs.

Accomplishments: The project simulated ice formation within million-molecule water droplets on six separate surfaces with differing contact angles under isothermal and cooling conditions (Figure 3.3). It has replicated GE's observed experimental results and deepened the understanding of freezing at the molecular level. To GE's knowledge, no other research group has executed such a comprehensive molecular dynamics study of droplets freezing on surfaces. Selecting the coarse-grained model of water interaction potential[**] for water proved to be especially effective, resulting in a 40× decrease in time to solution as compared to the extended simple-point charge interaction potential[††] originally envisioned for this project.

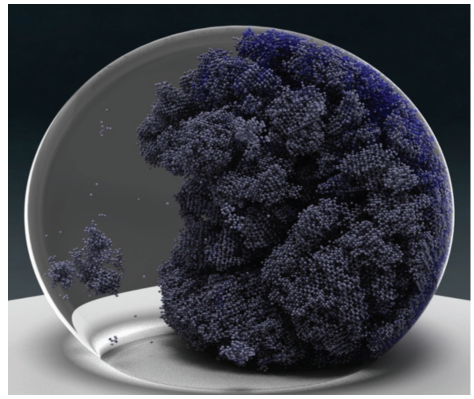

Figure 3.3. Snapshot from molecular dynamics simulation of ice crystal nucleation and growth. Location within the water droplet of ice nucleation varies depending upon temperature and surface contact angle. Visualization by M. Matheson (ORNL).

OLCF Contributions: OLCF computational scientist W.M. Brown and colleagues from the OLCF, Cray, and NVIDIA worked to prepare LAMMPS for use with GPUs within OLCF's CAAR (see Sect. 1.4.3). Brown and M. Yamada collaborated to take LAMMPS a step further and incorporated the new interaction for simulating water molecules. Choosing a new water model that is well suited for GPUs, skillfully recasting the model, and then running it on Titan's powerful hybrid architecture accelerated Yamada's freezing simulations. This achievement required many large-scale capability runs; 69% of the 38 million hours used was consumed at the capability scale (i.e., greater than 3,750 nodes). It also benefited strongly from the computational power availed through the use of Titan's GPUs, with 200 of the 350 total simulations being performed during the last months of the allocation, during which Titan's Kepler GPUs were available to the ALCC user program. In summary, by choosing a new water model that is well suited for GPUs, the time-to-solution decreased 40×. Skillfully implementing the model and then running it on Titan's powerful hybrid architecture accelerated Yamada's freezing simulations by an additional factor of 5×. Altogether, these improvements resulted in a 200× decrease in time-to-solution, and this speedup permitted a factor-of-eight increase in the number of runs executed over the number originally planned, tremendously increasing the new insights into how water freezes. The OLCF's Mike Matheson and Dave Pugmire provided visualization services.

Related Publication:

W. M. Brown and M. Yamada. "Implementing Molecular Dynamics on Hybrid High Performance Computers–Three-Body Potentials," *Comp. Phys. Comm.* 184 (12) 2785-–2793 (2013); doi: 10.1016/j.cpc.2013.08.002

Online Stories:

OLCF Staff Writer, "Titan Propels GE Wind Turbine Research into New Territory," *OLCF News* (October 25, 2013).

DOE, "Photo of the Week: Cold as Ice—Using Titan to Build More Efficient Wind Turbines," *Energy.gov Photo of the Week* (January 10, 2014).

HPC Wire, "Masako Yamada: Senior Scientist, GE Global Research," *HPC Wire People to Watch* (2014).

[**] V. Molinero and E. B. Moore, "Water Modeled as an Intermediate Element Between Carbon and Silicon," *J. Phys. Chem. B* 113, 4008 (2009).
[††] H. J. C. Berendsen, J. R. Grigera, and T. P. Straatsma, "The Missing Term in Effective Pair Potentials," *J. Phys. Chem.* 91, 6269 (1987).

GE Software, "Super Computing, Advanced Manufacturing, and the Industrial Internet," *Industrial Internet* (December 20, 2013).

OffshoreWIND.biz, "USA: GE Works on Solving Wind Turbine Ice Accumulation," *offshoreWIND.biz* (August 23, 2013).

The Daily Fusion, "GE Scientists Use Supercomputer to Solve Wind Turbine Ice Accumulation," *The Daily Fusion* (August 22, 2013).

Award Recognition:

GE Global Research received an International Data Corporation (IDC) HPC Innovation Excellence Award presented at the SC 2013 conference for this achievement.

International Data Corporation, "IDC Announces Winners of Sixth HPC Innovation Excellence Awards," *IDC Press Release* (November 19, 2013).

3.2.4 Rational Design of Efficient Organic Photovoltaic Materials: J-M Carrillo, W. M. Brown, ORNL, LAMMPS Early Science Project

Objective: Discover efficient designs for organic photovoltaic (OPV) solar cells, which are by their nature less efficient than (but also less expensive and more flexible than) inorganic photovoltaics.

Impact: This Early Science project performed the largest and longest molecular dynamics simulation of a promising OPV polymer system ever. For the first time, the simulation was performed for enough physical time to allow researchers to understand with a high degree of accuracy how the components of this system organize themselves and interact with one another. Understanding the formation, morphology, and domain size of bulk-heterojunction polymer blend active layers is critical for improving the performance of OPV solar cells (Figure 3.4).

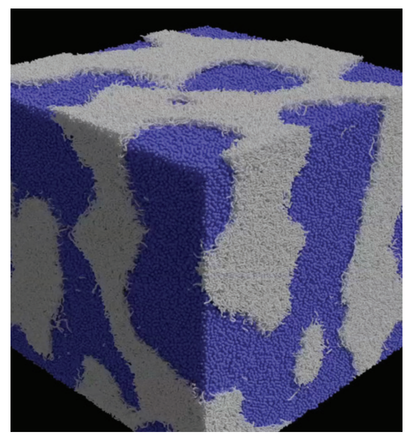

Figure 3.4. Snapshot of a coarse-grained molecular dynamics simulation of a phase-separated polymer blend (P3HT/PCBM in 1:1 weight ratio). The P3HT molecules are electron donors (white), and the PCBMs are electron acceptors (blue).

Accomplishments: Simulations of a polymer blend poly(3-hexylthiophene) (P3HT) and fullerene derivative [6,6]-phenyl-C61-butyric acid methyl ester (PCBM) determined both that you can mix these molecules and under what conditions. The project predicts both that increasing polymer chain length will decrease the size of the electron donor domains and that PCBM (fullerene) loading parameter results in an increasing, then decreasing, impact on P3HT domain size. These simulations on Titan are 27× larger and 10× longer in duration than the previous state-of-the-art published simulations. Calculations of this unprecedented size enabled this P3HT:PCBM bulk heterojunction formation to fully converge for the first time.

OLCF Contributions: OLCF computational scientist Mike Brown and other OLCF staff modified LAMMPS to efficiently take advantage of GPUs, thus delivering a speedup of 2.5 to 3 times over a comparable CPU-only system.

Related Publication:

J.-M. Carrillo et al., "New Insights into Dynamics and Morphology of P3HT:PCBM Active Layers in Bulk Heterojunctions," *Phys. Chem. Chem. Phys*. 15 (41), 17873 (2013).

Online Story:

OLCF Staff Writer, "Titan Sheds Light on Unknowns in Organic Photovoltaic Research," *OLCF News* (August 21, 2013).

Award Recognition:

GE Global Research received an IDC HPC Innovation Excellence award presented at the SC 2013 conference for this achievement.

International Data Corporation, "IDC Announces Winners of Sixth HPC Innovation Excellence Awards," *IDC Press Release* (November 19, 2013).

3.2.5 Titan Hosts Four of Six 2013 Gordon Bell Prize Competition Finalists

Four of six Gordon Bell Prize finalists used Titan to overcome complex computational challenges. From the dynamics of millions of crowded proteins to the billions of particles that radiate from passing plasma jets, ultrahigh-resolution simulations on Titan make for strong contenders.

Massimo Bernaschi, ICNR-IAC Rome, Biofluidic Systems

Researchers took advantage of Titan's power to simulate proteins in a realistic and, therefore, crowded environment. Such an approach is necessary because proteins in a living cell interact with other proteins and with surrounding fluids by changing their shapes, movements, and behaviors, sometimes dramatically. Their code, called MUPHY, for MUlti PHYsics simulator, studies the two-way interactions between proteins and fluids within a cell. The simulations used 18,000 of Titan's 18,688 nodes. It reached a sustained performance of 20 petaflops.

Online Story:

OLCF Staff Writer, "Peering into Cells One GPU at a Time," *OLCF News* (November 6, 2013).

Peter Staar, ETH Zurich, High-Temperature Superconductivity

Researchers scaled to the full Titan system, took advantage of the system's GPUs, and reached 15.4 petaflops, simulating high-temperature superconductors with the algorithm DCA+. The algorithm also made substantial headway with two nagging problems common to dynamic cluster quantum Monte Carlo simulations: the fermionic sign problem and the cluster shape dependency. In addition, the application took full advantage of Titan's energy efficiency. Simulation of the project's largest realistic clusters consumed 4,300 kilowatt-hours. The same simulation on a comparable CPU-only system, the Cray XE6, would have consumed nearly eight times as much energy.

The DCA+ team gained access to Titan in 2013 through an ALCC project, (PI: Thomas Maier, ORNL). Maier and collaborators were awarded 60 million Titan core-hours in 2014 through the INCITE program.

Online Stories:

L. Williams, "Superconductor Simulation Tops 15 Petaflops on Titan," *OLCF News* (November 12, 2013).

S. Ulmer, "Quantum Leap in Superconductor Simulation," *ETH Zürich News and Event*s (November 19, 2013).

Michael Bussmann, HZDR—Dresden, Plasma Physics

A team from Germany's HZDR–Dresden used Titan to simulate billions of particles in two passing cosmic jet streams. By modeling a well-known property of plasma turbulence called the relativistic

Kelvin-Helmholtz instability, which occurs where passing plasma jets collide, researchers were able to make out patterns of particle behavior—the inner workings of these faraway objects. Then they used radiative signatures to correlate plasma dynamics with radiation emitted during turbulence. The application, known as PICONGPU reached 7.2 petaflops on Titan, and gained access in 2013 through a DD project (PI: Michael Bussmann).

Online Stories:

OLCF Staff Writer, "Simulations of Plasma Turbulence Model the Inner Workings of Cosmic Phenomenon," *OLCF News* (November 11, 2013).

Tiffany Trader, "Cosmic Supercomputing Code Selected as Gordon Bell Finalist," *HPC Wire* (November 13, 2013).

Salman Habib, Argonne National Laboratory, Cosmology

Researchers simulated the evolution of trillions of interacting particles to simulate the evolution of the universe as it expands across billions of years. Using the Hardware/Hybrid Accelerated Cosmology Code a code called (HACC), they followed the formation of the structure of the universe, creating a comprehensive sky catalog to which scientists can compare instrumental observations. Notable for this achievement is that the HACC code, originally developed for the Road Runner hybrid accelerated machine at Los Alamos National Laboratory (LANL), was extensively developed by Habib and collaborators at Argonne National Laboratory (ANL), where they used Mira, the IBM Blue Gene Q (BG-Q) present at ANL, and Sequoia, the IBM BG-Q located at Lawrence Livermore National Laboratory (LLNL). As a finalist in the SC12 Gordon Bell Competition, using Sequoia and Mira, HACC attains unprecedented scalable performance—13.94 petaflops at 69.2% of peak and 90% parallel efficiency on 1,572,864 cores with an equal number of Message-Passing Interface (MPI) ranks, and a concurrency of 6.3 million.

The HACC team gained access to Titan in 2013 through an ALCC project (PI: Katrin Heitmann, ANL). In a tremendous demonstration of code performance portability, this team achieved full-machine-scale calculations on Titan, necessary for manuscript submission for the SC13 Gordon Bell competition, within 2 week of gaining user access in March 2013 prior to machine acceptance and transition to user operations. Habib and collaborators were awarded 100 million Titan core-hours in 2014 through the INCITE program.

Online Stories:

OLCF Staff Writer, "Code for Largest Cosmological Simulations Ever on GPUs Is Gordon Bell Finalist," *OLCF News* (November 5, 2013).

R. Brueckner, "Gordon Bell Finalist Uses Titan to Power Universe Sky Catalog," *Inside HPC* (November 6, 2013).

3.3 DIRECTOR'S DISCRETIONARY PROGRAM

2013 Operational Assessment Guidance

The Facility should describe how the Director's Reserve is allocated and list the awarded projects, showing the PI name, organization, hours awarded, and project title.

The OLCF allocates time on leadership resources primarily through the INCITE program and through the facility's DD program. The OLCF seeks to maximize scientific productivity via capability computing through both programs. Accordingly, a set of criteria are considered when making allocations, including the strategic impact of the expected scientific results and the degree to which awardees can make effective use of leadership resources. Further, through the ALCC program, the ASCR office allocates up to 30% of the facility's resources.

The goals of the DD program are threefold:

- preparation for leadership computing competitions (i.e., INCITE and ALCC) to improve and document application computational readiness;
- broaden the community of researchers capable of using leadership computing; and
- develop R&D partnerships, both internal and external to ORNL, to advance DOE and ORNL strategic agendas.

These goals are aligned particularly well with three of the four mission goals of the OLCF, namely

- to enable high-impact, grand-challenge science and engineering that could not otherwise be performed without to the leadership-class computational and data resources,
- to enable fundamentally new methods of scientific discovery by building stronger collaborations with experimental facilities as well as DOE offices that have large compute and data science challenges, and
- to educate and train the next-generation workforce grounded in the application of leadership computing to the most challenging scientific and engineering problems.

R&D partnerships are those aligned with DOE and ORNL strategic agendas. They may be entirely new areas with respect to HPC or ones in need of nurturing. Example projects are those associated with the ORNL Laboratory Directed Research and Development (LDRD) Program; programmatic science areas (fusion, materials, chemistry, climate, nuclear physics, nuclear engineering, and bioenergy science and technology); and key academic partnerships (e.g., the UT-ORNL Joint Institute for Computational Sciences). Examples of strategic partners in our DD program include the Center for Advanced Simulation of Light Water Reactors (CASL), Critical Materials Institute hub led by the Ames National Laboratory, Accelerated Climate Modeling for Energy (ACME) program, Fusion Simulation Project, Center for Nanophase Materials Sciences (CNMS), and large experimental facilities such as the Spallation Neutron Source and the ATLAS Experiment at CERN and Brookhaven National Laboratory (BNL). Also included in this broad category is the Industrial HPC Partnerships Program, providing opportunities for industrial researchers to access the leadership systems to carry out work that would not otherwise be possible. Through its Industrial HPC Partnerships Program, OLCF is achieving the original Congressional intend for the Leadership Computing Program by providing "Leadership Systems, on a competitive, merit-reviewed basis, access to researchers in United States industry, institutions of higher education, national laboratories, and other Federal agencies."[‡‡]

The DD program is also accessible to the general HPC community to carry out porting and development exercises for nascent and less-efficient applications. These performance enhancement projects range in scope from immediate INCITE preparation—designed to allow investigators the opportunity to test their codes' computational readiness on INCITE platforms—to somewhat longer-term projects involving improvement in algorithms and implementations. The ORAU-ORNL High-Performance Computing Grant Program provides modest grants of research funding, provided by ORAU, and modest grants of leadership computing time, provided by OLCF, to ORAU-member-university faculty members. The program was established in 2009 to encourage new and expand existing research

[‡‡] Department of Energy High-End Revitalization Act of 2004. Public Law 108–423—NOV. 30, 2004.

initiatives among ORAU member institutions using HPC systems. The program is a competitive grant program managed and funded by ORAU and open only to ORAU's member institutions.[§§]

The following are examples of DD program outcomes in expanding the leadership computing science community. Of the 139 DD projects operational at OLCF in 2013, 38 were in support of the development of proposals submitted to the 2014 INCITE Call for Proposals, and 9 of these proposals submitted were awarded allocations at OLCF. See Appendix D for a complete list of the 2013 DD projects.

The OLCF DD program also supports a variety of "data projects" that require data storage and bandwidth capabilities but few compute resources (see Section 4.2). Ongoing data projects include the Earth System Grid Federation, Data Sharing Project for the Center for Exascale Simulation of Combustion in Turbulence codesign project, and the Majorana Demonstrator Secondary Data Archive. In addition, infrastructure software such as frameworks, libraries, and application tools and support research areas for next-generation operating systems, performance tools, and debugging environments are often developed in DD projects.

The Resource Utilization Council makes the final decision on DD applications, using written reviews from subject matter experts. The actual DD project lifetime is specified upon award: allocations are typically for 1 year or less. The typical size of DD awards is roughly 3 million Titan core-hours but can range from tens of thousands to 20 million Titan hours or more.

Since its inception in 2006, the DD program has granted allocations in virtually all areas of science identified by DOE as strategic for the nation (Table 3.2).

Table 3.2. Director's Discretionary Program: Domain Allocation Distribution

Time Period	Biology	Chemistry	Computer Science	Earth Science	Engineering	Fusion	Materials Science	Nuclear Energy	Physics
2008	19%	8%	28%	4%	8%	15%	3%	1%	14%
2009	5%	3%	19%	6%	8%	6%	33%	1%	19%
2010	9%	6%	10%	8%	19%	6%	16%	3%	23%
2011	7%	1%	10%	19%	14%	0%	9%	13%	26%
2012	6%	1%	21%	14%	25%	5%	10%	1%	18%
2013	9%	4%	15%	15%	12%	8%	14%	4%	19%

Annual DD program utilizations are typically less than the allocable hours; that is, for example, all of the DD time is not allocated at the beginning of the calendar year. With this approach, the OLCF can remain flexible and responsive to new project requests and research opportunities that arise during the year. The leadership computing resources continue to be effectively used under this approach, as INCITE and ALCC users are not "cut off" when they overrun their allocation. Rather, they are allowed to continue running at lower priority to make use of potentially available time.

3.4 INDUSTRIAL HPC PARTNERSHIPS PROGRAM: ACCEL

The Industrial HPC Partnerships Program completed its fifth year in 2013 and adopted the moniker ACCEL: Accelerating Competitiveness through Computational Excellence. As a maturing program, ACCEL continues to help large and small companies supercharge their competitiveness and realize bottom-line benefits through access to the OLCF's leadership computing resources and computational expertise. Twenty-nine projects from 14 companies were under way during the year. These projects used more than 311 million Titan core-hours, representing approximately 11% of the total hours that Titan delivered in 2013. Of these projects, 15 were new in 2013, with companies receiving a total of 302 million Titan core-hours from the INCITE, ALCC, and the OLCF DD programs.

---

§§ ORAU/ORNL High Performance Computing (HPC) Grant Program.

Bosch, Caterpillar, KatRisk, and Saudi Aramco launched their first projects at the OLCF. Caterpillar and Saudi Aramco represent new industries for the ACCEL program (construction and mining equipment, oil and gas). And KatRisk is a startup, demonstrating that small companies also have challenging HPC requirements.

Six firms received DD allocations through a special ACCEL call for proposals to test out Titan's GPUs. This is discussed under the Innovation section (Section 4.1). Three companies successfully competed for larger-scale awards through ASCR's ALCC program: United Technologies Research Center, small business Ramgen Power Systems, and GE Global Research (two projects). Grants to these firms collectively represented approximately 30% of the total ALCC hours granted at the OLCF for the 2013–2014 program year. In addition, Bosch and Procter & Gamble (in partnership with Temple University) each received awards under the INCITE program.

This year ACCEL also helped launch two innovative, multiyear research collaborations with Ford and GM. New government regulations require that auto company fleets must average 54.5 mpg by 2025. GM and Ford are pursuing use of large-scale, predictive simulations to accelerate engine development to meet these looming efficiency and emissions goals while meeting customer demands for better performance at a lower cost. Their projects at the OLCF also are noteworthy in that project goals include optimizing two software applications important to the broader automotive industry and developing basic computational infrastructure that can be easily extended to engineering designs in other industries.

3.4.1 GM investigates fuel injectors

As part of GM's program to achieve greater automotive fuel efficiencies, it is converting more of its gasoline engine portfolio to so-called "spark-ignited direct-injection" (SIDI) technology and adopting multi-hole SIDI gasoline spray injectors. Better understanding of the gasoline spray formation and subsequent turbulent fuel-air mixing in the combustion chamber is critical. Multihole spray injectors offer great flexibility to control the orientation of spray and to engineer a variety of spray patterns. However, determining the optimum size, shape, and placement of the nozzle holes in order to produce the needed liquid or vapor spray patterns is a complex "trial and error" process. For example, a highly penetrating spray may lead to excessive liquid wetting that may result in higher hydrocarbon emissions. Also, the spray needs proper containment to achieve the correct mixture distribution to ensure smooth combustion.

GM turned to OLCF experts and Titan to help them better understand the effect of injector nozzle design on spray patterns and reduce the number of probable injector concepts and physical prototypes that must be built and tested. GM is using OpenFOAM, an open-source computational fluid dynamics (CFD) application that is becoming increasingly popular within the automotive, aerospace, and consumer products industries for larger-scale simulations. Two different versions of OpenFOAM were modified by OLCF researchers to compile on Titan and now are available on Titan for broader use. ORNL researcher Wael Elwasif (funded by the DOE Office of Energy Efficiency and Renewable Energy (EERE) Vehicle Technologies Program) then adapted an existing Python computational "wrapper" to enable GM to use OpenFOAM to conduct multiple simultaneous jobs to accelerate the process to reach an optimal design. This framework is open and also now available to any user at OLCF that wants to launch multiple simultaneous CFD jobs using OpenFOAM.

To date, several spray injector cases (with operating temperatures and pressure as the variables) were simulated and validated against experimental data from GM. Now high-fidelity simulations of these cases are being performed to assess grid refinement. The next steps will be to conduct a complete assessment of the sensitivity of the spray patterns to common adjustable design parameters, an important step toward GM's goal of creating an end-to-end assessment of injector design impact.

Access to Titan and OLCF expertise is enabling GM to demonstrate the utility of highly parallel simulations to accelerate and to optimize the design of these fuel injectors, and do so within the context of the overall performance of the engine.

3.4.2 Ford Seeks to Understand Cycle-to-Cycle Combustion Variation in Engines

Ford is using OLCF resources to pursue a different research path to achieve greater engine efficiency. Ford's project aims to gain a better understanding of the processes that create cycle-to-cycle combustion variation in a spark ignition engine when these engines run under so-called dilute conditions. Dilute combustion with cooled exhaust gas recirculation (EGR) is an effective technology for improving fuel efficiency while meeting future emissions requirements. Dilution with EGR allows decreased pumping and heat transfer losses, while permitting the continued use of very effective exhaust aftertreatment with a three-way catalyst.

Unfortunately, at higher EGR rates, cycle-to-cycle variation in combustion emerges as a significant limiting factor. Ford is partnering with OLCF experts and using Titan to develop a novel computational approach to address this challenge. Previous approaches required that many hundreds of engine cycles had to be computed serially (one cycle followed by another, similar to how the experiments are carried out) to reveal the critical statistics needed to better understand the extent of variations. Running that many cycles with a model detailed enough to capture the detailed in-cylinder combustion processes would be time prohibitive on Ford's in-house HPC system as it can take several months of wall-clock time. Through innovative algorithms developed by ORNL researchers, Ford is replacing current serial combustion cycle simulations with massively parallel, multiple simultaneous combustion cycle simulations.

The new algorithms were developed by ORNL researchers (Charles Finney, Miroslav Stoyanov, Dean Edwards, and Sreekanth Pannala funded by EERE Vehicle Technologies Program). They incorporated these algorithms into new software to sample both the operating parameters as well as the residual feedback effects and demonstrated that it could run on Titan to accelerate the computational study of the parameters contributing to this cycle-to-cycle variability. This software was then successfully coupled or "wrapped" with the Converge CFD code from small business Convergent Science. This wrapper enables Converge to use thousands or tens of thousands of processors to simultaneously initiate many cycle simulations with controlled perturbations in the initial conditions as well as feedback effects. Since Ford (and other automotive companies) relies on Converge to do its in-house in-cylinder combustion simulations, this wrapper was crucial so that Converge could take advantage of Titan's capabilities. This sampling wrapper is open and now available to OLCF users. It also can be used with any CFD software that is used for engine analysis and will help accelerate that software similarly.

Ford has used Converge with its new "wrapper" to test the viability of this massively parallel approach, and the simulation results are being processed. Ford started using Converge and its wrapper on its engine geometry where it already has data for validation. Once the simulations are validated, Ford would like to use them to understand the origins of cycle-to-cycle variabilities and how to modify design and operating parameters to eliminate them for next-generation, efficient automotive engines. Ford anticipates that it will reduce solution time for this problem from months to days or a week on Titan and transfer this capability to its in-house resources.

As ACCEL matures and helps industry solve more complex problems, it is helping DOE meet its strategic goal to maintain a vibrant US effort in science and engineering as a cornerstone of our economic prosperity with clear leadership in strategic areas.

Innovation

HIGH PERFORMANCE COMPUTING FACILITY
2013 OPERATIONAL ASSESSMENT
OAK RIDGE LEADERSHIP COMPUTING FACILITY

February 2014

4. INNOVATION

CHARGE QUESTION 4: Have innovations been implemented that have improved the facility's operations?

OLCF RESPONSE: Yes. The OLCF actively pursues innovations that can enhance facility operations. Through collaborations with users, other facilities, and vendors, many of these innovations are disseminated and adopted across the country.

4.1 INDUSTRIAL PARTNERING TO USE GPUS

The OLCF recognizes that many industrial HPC users may lack the experience with GPUs to efficiently run applications on Titan. To address this issue, the OLCF, through its industrial partnership program, Accelerating Competitiveness through Computational Excellence (ACCEL), invited select firms to submit proposals for a unique Titan access opportunity. This special call for proposals was launched to help industry expand its ability to harness the power of GPUs to accelerate problem solving. This opportunity provides both computing time on Titan and technical expertise, volunteered by Cray's Center of Excellence at the OLCF and by NVIDIA. In turn, the OLCF will learn from industry's questions and experiences about the impact of Titan's architecture on real-world industrial problems.

Seven ACCEL proposals were selected in 2013: two from Procter & Gamble and one each from Boeing, the United Technologies Research Center (collaborating with Pratt & Whitney), Caterpillar, General Electric, and Simulia (an independent vendor of Abaqus software). This is a very exciting set of projects from firms that are making a commitment to GPUs and that want access to Titan to explore scaling their codes and running test problems to see how they perform. Four of the projects launched in the late third and fourth quarters of 2013 and two are expected to launch early in the first quarter of 2014. The OLCF anticipates that this group of early industrial GPU users will have interesting results and "lessons learned" that can be shared with the scientific community and with other industrial users who may be considering adding GPU capability to their internal HPC environment. Each participant has committed to sharing information by publishing their results and by participating in conferences and workshops.

4.2 NEW WAYS TO RESPOND TO USERS' DATA NEEDS

As leadership systems grow in size, so potentially does the amount of critical data generated during and after the simulations. These data are, in essence, the intellectual capital of the research communities. The OLCF continues to focus attention on ways to facilitate scientific accomplishments through efforts to improve data generation, its movement, access, and analysis. The following represent activities initiated in CY2013 in response to user data needs.

- ***Data Liaisons***: At its inception the OLCF initiated the highly successful liaison support model, whereby a team of experts in computational and computing domains are employed by the center to provide technical, visualization, and scientific support to users. In 2013 the OLCF formalized the addition of data liaisons to its team. The data liaisons provide critical I/O and workflow assistance to enable users to complete the cycle of science discovery from simulation to data movement, storage, and analysis. Examples of their efforts are included in the science accomplishments described in Sections 3.2.1 and 3.2.2.

- ***Data Projects***: The OLCF's DD program awards compute time to projects that meet strategic goals of the center (e.g., INCITE preparation, code development) In 2013, in response to user requests for continued access to data following successful completion of projects, the OLCF initiated a beta-test of a new DD project type: the data project. Whereas projects awarded and tracked by the center focus on compute time, a select number of data projects were granted in 2013 in order to allow users to access data after the award of compute time was exhausted. The pilot includes common use cases: servicing these projects will allow the OLCF to develop a suite of quantifiable requirement metrics for data-centric projects, including measures such as storage capacity, bandwidth requirements, and latency requirements. These, in turn, directly map to necessary hardware and software infrastructure to support these types of projects.

Through the data project beta test and introduction of data liaisons, the OLCF seeks to understand the ways in which data creation, storage, curation, exploration, and technical support lead to scientific insight. Effective storage (both capacity and bandwidth) and curation are essential to allow "complete"—i.e. from creation to analysis to conclusion—use of the data to produce outcomes. Additional requirements, both in infrastructure and policy, come to the fore if those data are to be made available to a wider community of researchers. Importantly, the kind of requirement gathering undertaken by the assignment of data liaisons and the instantiation of data projects is very much in keeping with the successful methods the OLCF has used in the past to understand simulation requirements.

OLCF staff is at the forefront of activities to benefit the center and its users in generating and analyzing data, partnering with other organizations to develop next-generations tools and resources (see, e.g., Lustre activities in Section 4.3 and 4.4.1; digital object identifiers pilot in Section 4.5.1; and parallel data tool development in Section 4.8)

4.3 DEVELOPMENT OF THE SPIDER II FILE SYSTEM

In 2013, the OLCF successfully designed, acquired and deployed the Spider II parallel file system (PFS). Spider II is a centerwide Lustre PFS with a capacity of 32 petabytes and I/O bandwidth of more than 1 terabyte/second to OLCF users. Spider II comprises 36 Data Direct Networks (DDN) SFA12K40 cabinets, 288 I/O servers, and 20,160 2-terabyte Near-Line SAS disk drives. Compared with Spider I (the previous centerwide Lustre PFS at the OLCF), Spider II is a significant leap forward, providing a 4× speedup in transfer rate and a 3× increase in capacity while operating under the same power envelope and a smaller footprint. The increased capacity and performance will enable OLCF users to perform scientific simulations much more quickly.

The Spider II PFS was a complex engineering project. The project began with an analysis of production I/O characteristics of Spider I; the resulting baseline and benchmark suite were used to evaluate a broad range of existing and emerging storage and file system technologies. The OLCF's assessment of these technologies and recommendations for improvement were communicated back to respective vendors. The team then executed a full/open acquisition, using a highly detailed and flexible RFP that has become a model for similar storage acquisitions. The acquisition resulted in a subcontract award to DDN.

Staff completed the deployment and acceptance of the Spider II hardware in September; it is fully integrated in to the existing infrastructure. These activities included an upgrade of the OLCF's Scalable I/O Network from DDR InfiniBand to FDR. The Lustre router nodes on Titan were upgraded with the FDR cards. Mellanox device drivers were updated, and changes were made to the Cray and Lustre software stack to support the new architecture. These code fixes were provided back to their respective vendors for incorporation in to the public releases for those products.

Lustre is an open source project with developers from across the world and the OLCF team contributes significantly to its development. At the 2013 Lustre User Group conference, Intel, as one of the two principal maintainers of the code base, recognized both LLNL and ORNL for their contributions.

The team continues to perform periodic large-scale tests on Titan in order to fine-tune Spider II. Recent tests show that the OLCF was able to achieve 1.3 terabyte/second read and 1.2 terabyte/second write performance at the block level. At the file system level, the read I/O performance from Titan to Spider II with checksums turned off is currently just over 1 terabyte/second; with checksums turned on, the performance is at 894 gigabyte/second. We will continue to fine-tune these numbers through several optimizations.

4.4 TECHNOLOGY TRANSFER TO VENDORS

4.4.1 Lustre File System Fixes for Performance and Functionality

During the installation and testing of the 30 petabyte block storage solution, and during the introduction to production of the accompanying Lustre file system, the OLCF staff identified a series of critical performance and functionality issues among the separate components and systems. Many of these issues were related to a combination of new product releases and the large scale of the target solution. Problems were identified on the block storage controller, the disk drive device driver, the InfiniBand software stack, the motherboard for the OSSs, the LNET router, and the Lustre software stack. In each instance, OLCF staff was able to characterize the issue, provide a mechanism for reliably reproducing the issue, and then either fix it internally, or work directly with the vendor to identify the appropriate solution. These efforts led to greater functionality and performance on the block storage controller, a revised device driver for each of the Lustre file system's 20,160 drives, with better sequential read and write performance, and a high-performing FDR client for the Cray XK7 LNET routers. These improvements will better these products for all customers, not just ORNL. As a result of these efforts, the Lustre file system has met all performance goals, and been released to production for the Titan, Eos, and Rhea user communities.

4.4.2 TORQUE Resource Manager—Recognition of OLCF Contributions

The OLCF uses Adaptive Computing's TORQUE and Moab products for workload management. Don Maxwell, from the OLCF HPC Operations Group, was recognized by Adaptive for his contributions to the maturation of these products on the Cray platform since they were first ported to that platform at the OLCF roughly seven years ago. OLCF staff members made several significant contributions to that code base in the areas of security, scalability, and reliability. Most recently, a security issue was discovered at the OLCF that allowed a root exploit in TORQUE. By providing a command in the email portion of the qsub command, a user could run any command as the privileged root user. The OLCF provided a patch to Adaptive Computing to fix the issue, and all TORQUE sites, including most of the large Cray sites that use the TORQUE product, were notified to fix the issue. Over the past year, the OLCF also provided fixes to Adaptive Computing to address memory leaks and deadlocks found on Cray platforms, and a fix for phantom jobs that were stranded on the pbs_mom servers (*pbs_mom* is the process that starts a user's job script and ensures that it completes within its allotted time) that then prevented new jobs from starting.

4.5 INNOVATIVE DATA MANAGEMENT

4.5.1 Digital Object Identifiers Work Presented at DataCite Summit held at the National Academy of Sciences

Recent directives from the Office of Science and Technology Policy, the Office of Management and Budget (OMB), and the Office of the President outline a desire to provide free access to scientific data arising from taxpayer-funded research. Fulfilling this desire presents several interesting challenges in terms of data sharing, long-term data preservation, and stewardship. These issues suggest the need for a rich set of data services that can help shape and implement sophisticated data policies. For example, the need arises for a mechanism to uniquely identify and provide access to all supercomputer-produced data of interest. This mechanism would ideally support open sharing of data by identifying and curating data products of value, allowing authors to provide arbitrary annotations (metadata), and identifying major results stemming from projects.

In response to the new data sharing requirements, the OLCF has initiated a feasibility study to assess the promise of digital object identifiers (DOIs) and their potential use within an OLCF context. A DOI offers a unique, permanent identifier to help track data sets that are produced by users. Metadata about the data set is stored in along with the DOI name; metadata may include information such as a location, a description of the data set, and the processes used to produce the data set. DOIs provide a key first step toward establishing the desired data-sharing ability for the scientific community. OLCF conducted a survey of several DOI registration authorities [e.g., the DOE Office of Science and Technical Information (OSTI), EZID, ORCID] and picked OSTI due to its ease of use and flexibility and because it is freely available for DOE projects. OLCF has built a prototype workflow infrastructure that includes the following software elements:

- the ability for a user to request a DOI for a data set,
- the composition of the needed metadata into a DOI request/registration to OSTI,
- obtaining a DOI handle and logging it in a catalog for search and retrieval,
- the backend infrastructure that is necessary to archive the data into long-term storage (i.e., from user-managed space into a center-managed area),
- annotating that location in the catalog, and
- a process for retrieving the data set based on its DOI.

Part and parcel of obtaining a DOI is the implicit assumption that the data will be maintained for a long duration, requiring it to be moved to a permanent location. OLCF has further engaged in conversations with several of user communities to understand how DOIs can benefit their processes and to determine the granularity at which a DOI should be requested (e.g., collections of data sets).

DOIs also provide several other benefits. A DOI helps with data sharing even before the publication of a scientific paper. It allows us to elevate the importance of data products, treating them as citable entities much like publications. DOIs can also help facilitate more fine-grained reporting to sponsors and a new metric to quantify progress. Further, DOIs can help both users and the center help identify key data sets for long-term preservation, instead of indiscriminately archiving all data products. A DOI-based archival system can have the added benefit of making data sizes for manageable. Initial efforts at the OLCF have been well received by the user communities. OSTI invited OLCF to give a presentation on the use of DOIs for supercomputing at the DataCite Summit held at the National Academy of Sciences in September 2013 (authors: Terry Jones, Sudharshan Vazhkudai, and Doug Fuller). The feasibility study is being transitioned into a larger pilot project so that OLCF can move toward formally obtaining DOIs for OLCF data.

4.5.2 Integrity Crawler Tool for Data Quality Assurance

OLCF's High-Performance Storage System (HPSS) installation holds more than 35 petabytes of scientific data, accumulated over a period of 20 years. The data's value is incalculable as intellectual capital for the scientific community. It includes the output of simulations run across a broad range of computer architectures and scientific disciplines during the course of 2 decades of research. As computer architectures and software evolve, much of the data stored in HPSS cannot be reproduced.

Ensuring that this legacy of scientific data is conserved effectively and remains reliably available, which depends on the reliability and robustness of the HPSS and its components, is a crucial interest for OLCF. OLCF staff began work on the HPSS Integrity Crawler in September 2013. Its purpose is to catch problems early, thereby avoiding data loss. Development toward production deployment is ongoing and should be completed in early 2014.

The HPSS Integrity Crawler regularly performs quality assurance on archived data. It also monitors HPSS operations and reports to the HPSS administrators potentially troublesome conditions as they arise. The HPSS Integrity Crawler supports the following features.

- Architected as a plug-in framework, the crawler has the ability to coordinate a collection of modules, each focused on a specific operational aspect of the HPSS. The current implementation contains several pluggable modules, including the Checksum Verifier to validate file contents and the Tape Copy Checker to verify that the number of copies stored for each file is correct. Additional plug-ins, including the Drill Instructor plug-in to exercise the various components of HPSS to provide early detection of faults due to hardware, network, or other failure, and the Migration/Purge Record Ager to ensure that outstanding migration and purge requests are handled in a timely way, are planned.

- The crawler builds a sample space of files to examine during each run. The metadata collected during each run is cataloged with the expectation that, over time, the center will have aggregated a wealth of integrity information about its data holdings. The HPSS archive has tens of millions of files. Therefore, the sample space needs to be representative of the archive but small enough to make progress during each run. In compiling the sample space, the HPSS Integrity Crawler pays attention to a variety of attributes, including projects, size, file age, tape age, and class of service.

- A lightweight design that will not hinder regular HPSS operations.

- Information that is already available, such as checksums for files that have already been computed, is leveraged as part of the Crawler development to take advantage of all available resources.

4.5.3 Specializing in Data Transfer Use Cases

The Data Transfer Node (DTN) environment is a collection of 16 Linux nodes designed specifically to facilitate moving data between OLCF storage resources and from other laboratories and institutions. The nodes are connected to both the centerwide InfiniBand fabric to access Spider II, the OLCF centerwide Lustre file system. Each node has a 10Gb/s Ethernet connection to facilitate WAN transfers and transfers to and from the HPSS. They are configured and tuned for high-speed data transfers, and they provide users a dedicated resource to access and transfer data so that the workload does not disrupt OLCF compute resources. The nodes were crucial in the migration of data from Spider I, the former centerwide Lustre file system, to Spider II.

As the usage and demand for such a resource has grown, the OLCF has deployed different classifications of nodes designed to address more specific data-transfer use cases. Scheduled DTNs are managed by the Moab job scheduler to support workflows on Titan that have a data movement component. The OLCF recently deployed a distributed copy tool that leverages MPI running on the scheduled DTNs to stripe transfers across multiple processes and nodes to speed up large transfers. HPSS DTNs offload the data transfer operations from high-speed interconnect (HSI) commands initiated from OLCF compute platforms that utilize the HSI transfer agent, and then the nodes perform the transfers to and from the HPSS in a parallel fashion. Interactive DTNs allow users to manually perform transfers using tools such as gridftp, bbcp, and scp. These nodes also function as Globus Online endpoints, which has become a very popular method of data movement between sites, due to its ease of use. This method of deploying DTNs for specific use cases has allowed the OLCF to scale the number of nodes in a more targeted manner to meet user needs.

4.5.4 Improving Data I/O through Better File System Partition Selection

As storage systems scale, one of the biggest challenges is how to meet the I/O throughput requirements of scientific applications. In leadership systems, a number of applications run concurrently, sharing the parallel file system (PFS) and contending for storage resources, potentially resulting in variable application performance. To alleviate the I/O contention on the shared PFS (Spider II) and to guarantee I/O throughput, the OLCF identified the need to proactively schedule I/O based on user needs and resource availability. The center staff is addressing this situation by building a monitoring infrastructure and a partition selection tool that is based on file system load.

One of the critical components needed to build such I/O-aware smart tools is the ability to monitor the PFS activity at zero-overhead, meaning no impact to the user I/O traffic and no additional load on the storage while collecting the file system log data. OLCF has built a custom tool that logs file system activity from the RAID controllers on the out-of-band management network with negligible impact to real application workloads. The collected data provide the bandwidth and count of read and write operations over a period of 2 seconds. Staff also log the request size distribution for both read and write operations at every logical volume. The data are populated in a database, which then allows OLCF staff to perform useful analytics and to help build higher-level services.

Using the trace data, OLCF staff has developed a tool, fs-select, to select a file system partition based on the load. The Spider II file system has several partitions, and users may choose a partition without any knowledge of the file system workload. This can result in resource contention and load imbalance if multiple users pick the same partition. Based on the captured storage workload, fs_select is able to generate aggregate statistics, which are then used to suggest a specific file system partition to an application. Use of the selected partitions has the potential to alleviate I/O bottlenecks. fs_select has been integrated with the Torque/Moab workload manager so that an application is allowed to query the optimal partition at runtime.

4.5.5 Parallel Data Tools from Multi-Lab Collaboration

Current storage systems scale into the tens of petabytes of storage capacity and into the terabytes per second of data transfer performance. The OLCF's Spider II PFS is one example of this, with more than 30 PB's of capacity, and greater than 1TB/second performance. Leveraging the high performance available from massive-scale file systems such as Spider II requires scalable software. Current leadership-class applications make use of highly optimized I/O libraries and/or individually tuned I/O routines in order to realize the scalability inherent in these file system designs. However, user-level tools are still based on I/O primitives provided by single instances of an inherently serial operating system.

The parallel tools effort, a collaboration among ORNL, LANL, LLNL, and DataDirect Networks, is aimed at employing the scalability available from current supercomputers and parallel file systems to increase performance for everyday file and data management tasks. Current tools under development include a Lustre-aware parallel data copy tool, a parallel tape archive (tar) tool, a parallel filesystem scanning tool, and a parallel filesystem search tool. For familiarity, each tool supports the same user interface as its legacy Unix counterpart. The tools are implemented by libcircle, a common workflow parallelization library.

OLCF staff has already produced initial versions of the parallel tar, parallel filesystem scanning, and parallel search utilities; these tools have been tested in the OLCF environment, and additional development of these tools continues among the collaborative partners.

4.6 INNOVATIVE APPLICATION SUPPORT

4.6.1 Nurturing the Ecosystem for Hybrid Programming

One of the distinctive features of the Titan system is its hybrid environment, in which GPU accelerators are used to provide a significant portion of its computational capability. Although this approach has been successful in bringing a significant increase in capability to the OLCF in a cost- and power-efficient way, it also requires users to adapt to the new hybrid programming environment in order to make effective use of the system. To facilitate adaptation to hybrid computing, the OLCF offers several different programming environments, which can be used in applications individually or in combination.

CUDA. The first and lowest level of these is CUDA, which is a C-based language for GPU accelerators developed by NVIDIA and targeting NVIDIA accelerators (such as Titan's Kepler K20X processors). Being implemented by and tuned to NVIDIA accelerators, well-written CUDA is generally expected to provide the best possible performance on a system such as Titan; however, code portability is sacrificed, as CUDA only works on systems with NVIDIA GPUs.

OpenCL. OpenCL, like CUDA, is a low-level approach, although it is based on a set of APIs rather than on specific architectural implementations. Different from CUDA, OpenCL is an open standard with multiple implementations targeting many-core CPUs, GPU accelerators from different vendors and other accelerators, such as Intel's Xeon Phi. Benchmarking shows that current OpenCL implementations lag significantly behind CUDA implementations in performance, which has led to limited interest so far among OLCF users. However, interest is growing as code teams recognize the value of code portability.

OpenACC. OpenACC is a higher-level approach, using directives (specially structured comments) and APIs to control the offloading of certain computations to the accelerator. It provides higher levels of both code portability and performance portability than CUDA or OpenCL at some cost in absolute performance. It is an open standard with multiple implementations targeting a variety of accelerator hardware types.

OpenACC is also serving as a test bed for the slower-moving OpenMP standard. Although historically focused on threaded programming for many-core processors, OpenMP 4.0 (released in November 2013), with input from the OpenACC community, is also beginning to support accelerators. As a long-standing open standard, OpenMP tends to be relatively conservative, with minor new revisions of the standard released approximately every 2 years and major releases delivered every 4 years. On the other hand, OpenACC updates are issued twice per year. Over the last several years, the OLCF has placed a particular emphasis on OpenACC and the ecosystem surrounding it in order to ensure that its users and the broader HPC community have a robust and usable set of programming environments to choose from and the information and guidance to allow them to make informed choices about which approaches best suit their needs.

Recognizing the early proliferation of proprietary directive-based accelerator programming environments, the OLCF became one of the leading advocates for the formation of the OpenACC standards organization, launched in 2011. Subsequently, in 2012, ORNL became one of the first nonvendor organizations to join this organization. Since the inception of OpenACC, the OLCF has worked with the community and made strategic investments with vendor partners to ensure that OpenACC provides a viable and productive environment for OLCF users. The OLCF's commitment includes working within the standards body to ensure that the needs of OLCF applications are addressed, often by providing detailed motivating examples. This concept is exemplified by the face-to-face meeting of the OpenACC organization hosted at ORNL in August 2013. In addition to the usual standards meeting, the local hosts, ORNL's Oscar Hernandez and CAPS Enterprise's Jean-Charles Vasnier arranged a day's worth of "deep dives" with a variety of researchers presenting their applications' needs to the members of the standards body. In recognition of efforts like these, ORNL researcher Oscar Hernandez was appointed OpenACC's director of user adoption.

Through development contracts with vendor partners CAPS Enterprise (an OpenACC compiler vendor), Allinea (vendor for the DDT debugger) and TU-Dresden (originator of the Vampir profiling tool), the OLCF encourages prototyping and rapid implementation and release of key capabilities (of particular interest to OLCF users) preceding standardization (such as the OpenACC performance monitoring API) and immediately upon standardization.

Through its long-standing membership in the OpenMP Architecture Review Board, ORNL is also helping to ensure that features of OpenACC are appropriately incorporated into OpenMP as they mature.

4.6.2 Enabling Application-Level Resilience with a Fault-Tolerant Message-Passing Interface

Today's HPC applications rely primarily on checkpoint/restart for resilience in the face of system failures. Many applications could, in principle, use other approaches that might be more resource efficient, but there is little production-quality infrastructure to support them. For example, most implementations of the MPI standard will, by default, kill the entire job in response to the failure of a single node. Thus, even if a given application could in principle tolerate the node failure, the standard MPI infrastructure does not support such a possibility.

ORNL researchers, supported by the OLCF, have, for the last several years, played a significant role in the MPI community in helping to close this gap by working to define, implement, and demonstrate the proposed User-Level Fault Mitigation (ULFM) interface. As part of the MPI Forum's Fault Tolerance Working Group, ORNL researchers have been engaged in reviewing and evaluating HPC fault tolerance. In support of this effort, the ORNL team has developed STCI, a new low-level runtime infrastructure that provides the fundamental capabilities required to (1) allow an MPI job to "run through" a node failure without the entire job failing and (2) refine the application-facing interface that allows for the notification of node failures and recovery under control of the application. The team is working with OLCF users to demonstrate the use of ULFM in their applications, which the MPI Forum has considered an important prerequisite to acceptance of ULFM into the next MPI standard.

4.6.3 Improving Performance by Reviewing aprun Usage

The OLCF has developed aprun-usage, a new module that reviews a user's aprun flags while a job is running. The module provides feedback that could lead to improved performance by reducing wall time, run-time variability, or both.

Titan's AMD Opteron CPUs feature 16 cores but only eight floating-point units (FPUs). By default, aprun will place 16 processes per node. Applications that are primarily floating-point intensive may have pairs of processes contending for the shared FPU.

Some users will naively restrict their applications to use eight or fewer processes per node in an attempt to avoid FPU contention. However, this strategy requires twice as many nodes for a given job size and may still not avoid the problem. To effectively avoid the contention, they must pass the appropriate flags to distribute the processes (e.g. -j 1). If they do not use the correct flags, aprun will place them on the first N cores, and they will still contend over the FPUs even though they are using eight or fewer processes per node. If it is properly spaced, the job will run 1.4 to 1.7 times faster (i.e., wall time will be reduced by 30% to 40%). The speedup is typically less than 2.0 due to periods when the application does not use the FPUs, such as during I/O.

The aprun-usage module detects aprun calls that fail to properly place their processes and prints an advisory with a URL to an NCCS webpage with details on how to improve the performance. The module also suggests that the user turn on core specialization, when appropriate, to reduce jitter (i.e., runtime variability due to system noise). Lastly, the module also warns if the user specified more OpenMP threads than the user requests per process from aprun.

This module was developed in 2013 and deployed in Q4 as an optional module. Beginning in late January 2014, it will be loaded for all users by default. Over a 30 day period in Q4 2013, users submitted 29,784 jobs. Of these, the aprun-usage model detected 712 jobs that were not using proper placement to avoid FPU contention.

4.7 INNOVATIVE HARDWARE MONITORING AND DIAGNOSTICS

4.7.1 System Performance Tracking Tools

As part of the root cause analysis of early performance and failure issues on the Cray XK7, the OLCF developed a series of diagnostic tools that could identify underperforming or failing hardware. These tools were frequently used to screen the system after a reboot before allowing production jobs to start. While High Performance Linpack (HPL) eventually became the primary diagnostic tool, others were used in the discovery phase to identify slow performance, PCI lane degrades, and similar issues. These tools continue to be used today during screening procedures.

Another diagnostic tool was developed that could isolate individual GPUs that were contributing to an incorrect HPL result or residual. Because this fault was intermittent, it typically could be driven out only with large-scale or full-system HPL runs. However, due to the inherent method for calculating the HPL result and residual, the individual node that was causing the error could not be identified. The OLCF developed a method of isolating the offending hardware, using the concept of a virtual cabinet, placing a potentially faulty node in a virtual cabinet of known good nodes and executing a series of HPL runs among the members of the virtual cabinet to identify the failing node.

4.7.2 TopoBW MPI Program for Benchmarking and Diagnostics

The OLCF developed a benchmark and diagnostic application called TopoBW, a user-level MPI program that stresses the Cray Gemini network by saturating each link on the system and measuring the resulting performance. Using topology-aware algorithms, it can detect previously undiscovered performance problems. While Cray provides tools for assessing the health of the Gemini network, those diagnostics can only be run off line while the system is down for maintenance. TopoBW can be run while the machine is in production, and can even be run on a contiguous subset of the machine. TopoBW supersedes the capabilities of the offline tool as well, demonstrating an ability to identify failing connections that the offline diagnostics cannot.

4.7.3 Liebert XDP Protocol Translation Rolled into Centerwide Monitoring

The OLCF investigated and implemented a system to translate the native protocol used by the Liebert XDP units used to cool the Cray XE/XK models into a protocol that could be interpreted by the centerwide monitoring system. The Liebert XDP units use a data communication protocol for building automation and control networks called BACnet. This is a global standard used in the refrigeration, air-conditioning, and heating industries, but with little or no adoption by available open source monitoring solutions. The open source monitoring solutions primarily depend on the Simple Network Management Protocol (SNMP) for communicating with devices; the OLCF implements its monitoring in the same manner. The ready solution to the integration of the Liebert XDP units in to the existing monitoring solution was to find a way to translate BACnet messages into SNMP. A device developed by Chipkin Automation Systems provided the needed protocol translation. Once that device was configured and deployed, custom SNMP scripts and queries were developed, and now the cooling system data are presented to the OLCF's Nagios server to provide device status, alarm notification, threshold errors, and similar information. With access to this data, graphs were also generated to look for outliers and to trend the data. Although Nagios can be used to communicate specific thresholds, alarms, and status, problems such as hot spots, high humidity, high chilled water temperatures in a portion of a computer room, or an underperforming cooling unit are more easily found by aggregating similar data series in one graph. This graphing capability was implemented using Multi-Router Traffic Grapher and is now incorporated in to the center's monitoring efforts.

Risk Management

HIGH PERFORMANCE COMPUTING FACILITY
2013 OPERATIONAL ASSESSMENT
OAK RIDGE LEADERSHIP COMPUTING FACILITY

February 2014

5. RISK MANAGEMENT

CHARGE QUESTION 5: Is the facility effectively managing risk?

OLCF RESPONSE: Yes, The OLCF has a very successful history of anticipating, analyzing and rating, and retiring both project- and operations-based risks. Our risk management approach uses the Project Management Institute's best practices as a model. Risks are tracked and, when appropriate, retired, recharacterized, or mitigated. A change history is maintained for historical reference.

The major risks currently being tracked are listed and described below. Any mitigations planned for or implemented are included in the descriptions. As of this writing, the OLCF has one "high" operational risk: that the facility will continue having difficulty finding sufficient staff. To address the risk, staff recruiting and retention efforts have received increased emphasis.

5.1 RISK MANAGEMENT SUMMARY

The OLCF's Risk Management Plan (RMP) describes a regular, rigorous, proactive, and highly successful review process. The RMP is reviewed at least annually and is updated when necessary. The plan covers both OLCF operations and its various projects. Each project execution plan refers to the main RMP but may incorporate some tailoring specific to the project. Risks are tracked in a risk registry database application that is capable of tracking individual project risks separately from operations risks.

Operations and project meetings are held weekly, and risk, which is continually being assessed and monitored, is usually discussed at the meetings. At least monthly, specific risk meetings are held, attended by the federal project director, facility management, OLCF group leaders, and others as required. When assessing risks, the OLCF management team focuses its attention on the high and moderate risks as well as any low risks within the impact horizons associated with the risk. Trigger conditions are stated in the "Risk Notes" narrative section of the register when appropriate. Early and late risk impact dates are recorded as well. Risk owners are expected to be proactive in tracking any trigger conditions and the impact horizons of the risks for which they are responsible and to bring appropriate attention to management of those risks, whatever the risk-rating level.

The OLCF reports current high- and medium-level risks to the DOE program office as part of its monthly operations report. At the time of this writing, 21 active entries are in the OLCF operations risk register. They fall into two general categories: risks for the entire facility and risks particular to some aspect of it. Across-the-board risks are concerned with such things as safety, funding/expenses, and staffing. More focused risks are concerned with reliability, availability, and use of the system or its components (e.g., the computing platforms, power and cooling, storage, networks, software, and user interaction). Forty-five risks are tracked for the OLCF-4 project that is active at the time of this report. A project risk may be listed below if it could also significantly affect operations.

Costs for handling risks are integrated within the budgeting exercises for the entire facility. Risk mitigation costs are estimated as any other effort cost or expense would be. For projects, a more formal bottom-up cost analysis is performed on the work breakdown structure. However, for operations, costs of accepted and residual risks are estimated by expert opinion and are accommodated as much as possible in management reserves. This reserve is continually reevaluated throughout the year.

5.2 CURRENT RISK STATUS

The scope of operations risks remains relatively stable from year to year. Adequate funding is always a concern, and OLCF's mission of continual innovation requires both scientists and OLCF staff to make frequent adjustments to accommodate new technologies. Recently, OLCF has experienced some staffing difficulties because of a highly competitive job market.

The risks listed in the following tables may appear in more than one table. The table column labeled "Section" indicates where the detailed description of each listed risk can be found in Section 5.8.

5.3 MAJOR RISKS TRACKED IN THE CURRENT YEAR (2013)

Major CY2013 Risk	Section	Rating	Notes
1006—Inability to acquire sufficient staff	5.8.1	High	OLCF is having difficulty acquiring adequate qualified staff because of a highly competitive job market.
1101—Insufficient funding to meet DOE commitments (FY2014)	5.8.2	High (recently retired)	Uncertainty is a concern. Annual budgets are set with guidance from the ASCR office, but actual allocated funds are unknown until Congress passes funding bills. Continuing resolutions are common, and the center often goes several months before actual funding is resolved. As of this writing, adequate funding for FY2014 has been appropriated, thus retiring this risk in early 2014.
361—Scientists decline to port to heterogeneous architecture	5.8.3	Medium	Porting is difficult. Mitigation includes in-house experience and staff and user training program development. The trends are good because the 29 projects awarded 2014 INCITE allocations at the OLCF all either have GPU versions of their codes or strong plans to use GPUs within 2014. Even so, this remains a long-term concern.
906—Programming environment tools may be insufficient	5.8.4	Medium	Mitigation includes subcontracts with tool vendors and the creation of a special OLCF Software Tools team.
917—Robust support will not be available to ensure portability of restructured applications	5.8.5	Medium	Remains a concern.
948—Lack of infrastructure for an exascale system	5.8.6	Medium	Long lead time will be required to resolve. Remains a concern.

5.4 RISKS THAT OCCURRED DURING THE CURRENT YEAR AND THE EFFECTIVENESS OF THEIR MITIGATIONS

Risk that Occurred	Section	Rating	Notes
124—Storage system reliability and performance problems	5.8.7	Medium	The center worked with the vendor to redesign and replace the motherboards on the object storage servers to correct a design flaw in FY2013. Risk remains active because it continues to be a concern.

Risk that Occurred	Section	Rating	Notes
412—Inadequate system availability.	5.8.8	Low	Manufacturing defects in the Titan upgrade (OLCF-3) hardware required the rework of every blade in all 200 cabinets. This required removing portions of the system from production so that the rework could be performed. Performance targets were adversely affected. After the rework, acceptance of the system was completed and Titan was put into production. Continued monitoring of node failures showed that while within predicted specifications, there was a disturbing trend in the number of a specific type of failure. Ultimately, the OLCF, Cray, and ASCR agreed that a preventive maintenance period to do further board work was appropriate. The OLCF was able to work with our users to allow us to meet our 2013 operational metrics.

5.5 RISKS RETIRED DURING THE CURRENT YEAR

Retired Risk	Section	Rating	Notes
979—Insufficient funding to meet DOE commitments (FY2013)	5.8.9	Medium	The risk did not occur. Funding was sufficient for FY2013.
1001—Leadership computing performance target is not achieved (CY2013)	5.8.10	Medium	Risk did not occur. Continued improvement in application readiness by the OLCF Scientific Computing Group helped achieve targets, as did establishing job-queue policies with high preference for leadership jobs and OLCF's continued involvement with the INCITE proposal-selection process such that leadership-class projects received preference.
913—Long-term support for Lustre	5.8.11	Low	Long-term support contract in place. Intel has acquired Whamcloud, and Lustre is thriving in this setup. Further, OpenSFS is ensuring the continued progress of Lustre.
994—Costs for site prep for the disk acquisition may come in higher than estimated in original project budget.	5.8.12	Low	Risk did not occur.
1009—Impact of new travel regulations on the OLCF Training Program.	5.8.13	Low	The risk did not occur. The center was able to mitigate the risk through webcasting so that anyone who wanted to participate in OLCF workshops could do so free of charge. The center had the same number of or more participants at its workshops this year between on-site and webcast attendees.

5.6 MAJOR NEW OR RECHARACTERIZED RISKS SINCE LAST REVIEW

New Risk	Section	Rating	Notes
1006—Inability to acquire sufficient staff	5.8.1	High	Recently added. OLCF is having difficulty acquiring adequate qualified staff because of a highly competitive job market.

New Risk	Section	Rating	Notes
1101—Insufficient funding to meet DOE commitments (FY2014)	5.8.2	High (recently retired)	Uncertainty is a concern. Annual budgets are set with guidance from the ASCR office, but actual allocated funds are unknown until Congress passes funding bills. Continuing resolutions are common, and often several months will pass before actual funding is resolved. As of this writing, adequate funding for FY2014 has been appropriated, thus retiring this risk in early 2014.
721—Lustre metadata performance continues to impact applications	5.8.14	Medium	Mitigation includes participation in OpenSFS to develop features to improve the metadata server performance, deploying multiple file systems, and deploying the Distributed Namespace feature (DNE) in Lustre. The probability of the original appreciation of the risk was changed from Medium to High and the scope impact was changed from Medium to Low. These changes did not affect the overall original risk rating of Medium. The overall residual risk rating of Low did not change because it is believed that the mitigation efforts will be effective. Remains a concern.

5.7 MAJOR RISKS FOR NEXT YEAR

The major risks for next year will be similar to the major risks tracked this year (see the tables in Sections 5.3 and 5.4). However, as Titan has now completed acceptance and users are effectively using the system, many of the risks of such a new architecture are lower than the level at which the center had rated them last year.

5.8 DETAILED RISK DESCRIPTIONS

5.8.1 ID# 1006—Inability to Acquire Sufficient Staff

Risk Owner	Arthur S. Bland, OLCF Project Director		
Probability	High		
Impact	***Cost:*** Low	***Schedule:*** Low	***Scope/Tech:*** Low
Rating	HIGH		
Status	Accepting the risk		

The OLCF has difficulty acquiring adequate qualified staff because of a highly competitive job market. The risk is that desired work outcomes will not be achieved; some important tasks may be postponed or eliminated, and/or more current staff will become dissatisfied from overwork or missed opportunities to work on preferred assignments. The effect could be missed performance metrics, user dissatisfaction, or increased staff dissatisfaction.

Trigger: Open positions >10% of available positions

Although the cost, schedule, and technical impact ratings are all low, the risk is rated High because of "Other" impacts, such as those to OLCF's or ORNL's reputation as a preferred place to work.

The OLCF has increased its emphasis on both recruitment of new staff and retention of existing staff. Should management become aware that work outcomes might be impaired, temporary help may be obtained from other ORNL resources or contracts may be sought with external sources.

5.8.2 ID# 1101—Insufficient Funding to Meet DOE Commitments (FY2014)

Risk Owner	Arthur S. Bland, OLCF Project Director		
Probability	Low		
Impact	***Cost:*** High	***Schedule:*** Medium	***Scope/Tech:*** High
Rating	MEDIUM		
Status	Retired: The risk did not occur. Adequate funding for FY2014 has been provided.		

Annual budgets are set with guidance from the ASCR office, but allocated funds are unknown until Congress passes funding bills. Continuing resolutions are common, and often several months will pass before actual funding is resolved. The risk is that the center may have to delay some purchases, activities, hiring, etc., or adjust lease payment schedules, resulting in high costs or schedule delays.

Trigger: Intelligence on congressional or DOE funding capabilities and priorities

The center will maintain close contact with the federal project director and ASCR Program Office to understand the changing funding projections so that alternate plans can be made in sufficient time. Where possible, the center will structure contracts to accommodate flexible payment terms. In January 2014, Congress passed and the president signed an appropriation bill meeting the budget requirements of the OLCF, thus retiring this risk.

5.8.3 ID# 361—Scientists Decline to Port to Heterogeneous Architecture

Risk Owner	Jack C. Wells, NCCS Director of Science		
Probability	Medium		
Impact	***Cost:*** Low	***Schedule:*** Medium	***Scope/Tech:*** Low
Rating	MEDIUM		
Status	Mitigating the risk		

Common to all programming models is the need to structure and/or restructure codes to express increased hierarchical parallelism on today’s hybrid multicore architectures. This is necessary on all high-performance architectures to achieve good performance. Beyond this restructuring, one needs to use relatively new programming models to “offload” the computation to the GPU in GPU-accelerated hybrid architectures. The risk is that some users will decline to port or will delay porting of their applications to this new architecture because of the difficulty or cost. As a result the OLCF would expect to see a decrease in the number and/or quality of proposals submitted to allocation programs such as INCITE and ALCC.

Trigger: A decrease in the number and/or quality of proposals submitted to allocation programs such as INCITE

The original risk evaluation rated this risk as High. Mitigation with outreach, training, and the availability of libraries and development tools have ameliorated some initial user resistance. The marked improvement of compiler directive technology from Cray, CAPS, and PGI (including the OpenACC standardization) is overcoming some technical barriers for computational scientists to port and achieve acceptable performance running on hybrid, accelerated architectures. Additionally, the Tools team is leveraging LDRD and other investments to develop tools to assist users in porting their codes. Of the 29 proposals awarded INCITE projects at OLCF for 2014, 16 had a computational readiness score of greater or equal to 4 out of 5. Many applications teams appear to be porting their codes.

5.8.4 ID# 906—Programming Environment Tools May Be Insufficient

Risk Owner	David E. Bernholdt, Group Leader, Computer Science Research		
Probability	Medium		
Impact	***Cost:*** Medium	***Schedule:*** Low	***Scope/Tech:*** Medium
Rating	MEDIUM		
Status	Mitigating the risk		

The OLCF-3 system (Titan) relies on GPU accelerators for the bulk of its computational capability. The programming environment for OLCF-3 may not provide users with tools with which they are familiar, comfortable, and experienced and may not offer the levels of performance expected on the new system. If the programming environment is not productive for the users, they may withdraw from using the OLCF in favor of other centers.

Trigger: Concerns reported by user-application liaisons

The center created a Software Tools Group within the NCCS to own the problem. The center surveyed users on their requirements in this area and on the adequacy of the tools available or planned. It found that most of the primary tools from the OLCF-2 environment had plans to extend useful functionality for the OLCF-3 system. Where it found gaps, the center initiated contracts with vendors to accelerate their development and to add key functionality needed for the OLCF-3 system. These activities were moved into the OLCF-3 project as an initial risk mitigation. The center monitored the progress of the tools developers and checked out early versions of the tools on new Fermi processors in Jaguar and on other GPU-enabled systems to ensure the compatibility with existing programming models. The center has developed portable programming models (through our vendor partners) such as the directive-based OpenACC standard and the OpenMP directives for accelerators. Today, the center is a member of the OpenMP and OpenACC standards committees to push for needed improvements and eventual consolidation of these programming standards. It is also offering training to its users in how to use the programming models as well as the programming tools. The center has also contracted with Allinea, Dresden, and CAPS to have on-site user support to assist our users with the tools.

5.8.5 ID# 917—Robust Support Will Not Be Available to Ensure Portability of Restructured Applications

Risk Owner	Bronson Messer, Scientific Computing		
Probability	Medium		
Impact	***Cost:*** Medium	***Schedule:*** Low	***Scope/Tech:*** Medium
Rating	MEDIUM		
Status	Mitigating the risk		

The programming model that the center proposes requires a restructuring to utilize the standard distributed memory technologies in use today (e.g., MPI, Global Arrays) and then a thread-based model (e.g., OpenMP or Pthreads) on the node that captures larger-granularity work than is typically done in current applications. In the case of OpenMP, the compiler can facilitate and optimize this thread level of concurrency. This restructuring is agnostic to the particular multicore architecture and is required to expose more concurrency in the algorithmic space. Our experience to date shows that the center almost always enhances the performance with this kind of restructuring. The use of directives-based methods will allow the lowest level of concurrency to concomitantly be exposed (e.g., vector- or streaming-level programming). This means that the bottom level of concurrency can be directly generated by a compiler. The center expects that this kind of restructuring will work effectively with portable performance on relevant near-term architectures (e.g., IBM BG/Q, Cray Hybrid, and general GPU-based commodity cluster installations). However, restructured applications will be able to make use of several programming models—CUDA, OpenCL, OpenACC, or even parallel thread execution and other library-based approaches (e.g., OLCF's Geryon)—to expose the lowest (vector-like) level of concurrency.

The risk is that robust versions of OpenACC will not be available for other contemporary platforms. Also, OpenCL could be lacking on OLCF-3's platform and that OpenCL would remain lacking on the Titan platform.

The effect would be that applications run on Titan could be developmental "dead ends," due to poor performance, lack of a full set of features, or other problems. Users will have to work around these issues or change programming models.

Trigger: Intelligence on deficiencies in support applications

Multiple instantiations of compiler infrastructure tools will be adopted to maximize the exposure of multiple levels of concurrency in the applications. This will be abetted by publishing the case studies and experience gained from working with the six OLCF-3 project applications coupled with the appropriate training of our user community. The OLCF will work with vendors to continue to improve compiler technology and other tools. Additionally, the center has worked with compiler vendors to help form and promote OpenACC as a new standard aimed at providing a portable way to program for accelerator-based systems that is transparent to nonaccelerator systems.

5.8.6 ID# 948—Lack of Infrastructure for an Exascale System

Risk Owner	James H. Rogers, NCCS Director of Operations		
Probability	Low		
Impact	***Cost:*** High	***Schedule:*** High	***Scope/Tech:*** Medium
Rating	MEDIUM		
Status	Accepting the risk		

DOE's long-term plans include pre-exascale and exascale systems before the end of this decade. ORNL has a plan to provide the space, power, and cooling to support these goals, but there is risk that the systems will be significantly larger or use more power than projected and that the planned facilities will be insufficient. The risk is that exascale will not be achieved in a timely and effective manner with the effect of DOE mission goals not being accomplished at ORNL.

Trigger: Intelligence on the size and power requirements of proposed systems

ORNL has a plan to house the exascale system in building 5600 by moving other systems out of the building. However, the much-preferred approach would be to construct a new building that is designed for exascale from the beginning. OMB has rejected third-party financing as a method of building such a facility so this will need congressional line-item funding. There is an additional "Other" impact rating of High to reflect concerns of not meeting DOE mission goals adequately.

5.8.7 ID# 124—Storage system reliability and performance problems

Risk Owner	Sudharshan S. Vazhkudai, Group Leader, Technology Integration Group		
Probability	Medium		
Impact	***Cost:*** Low	***Schedule:*** Low	***Scope/Tech:*** Medium
Rating	MEDIUM		
Status	Accepting the risk		

Risk narratives were updated to reflect recent problem workarounds conducted in the summer of CY2013 to repair storage system motherboards.

Hardware or software bugs can cause the storage system to exhibit reduced reliability and performance. The storage system upgrade could exhibit reliability and/or performance issues during operations. An unreliable, underperforming storage system can result in intermittent failure, and unreliable performance or underperforming.

Trigger: Additional problems encountered during the first few month of the OLCF-3 file system being put into full production.

The storage system seems to be operating well now, but the risk will remain active for a while longer into CY2014 during continued observation.

5.8.8 ID# 412—Inadequate System Availability

Risk Owner	Kevin G. Thach, Group Leader, High-Performance Computing Operations		
Probability	Low		
Impact	***Cost:*** Low	***Schedule:*** Low	***Scope/Tech:*** Low
Rating	LOW		
Status	Accepting the risk		

This risk was downgraded to Low at the end of CY2013.

Availability and stability of systems is critical to users. There is a risk that the system stability and availability may not be sufficient to meet user needs or our DOE operational metrics. Projected FIT rates of the Kepler chip are worse than the FIT rates for CPUs; there is a risk that the system may not be stable enough to meet these operational requirements. This could mean loss of productivity, missed project deadlines, and user dissatisfaction.

Trigger: Measured trends of stability

The center will continue existing policies that control availability: minimize maintenance downtimes, coordinate upgrades, maximize fault-tolerant hardware and software, etc. It will measure availability and stability and use those results to detect trends in time to take remedial action. Working closely with NVIDIA and Cray, the center will characterize failures and develop responses should FIT rates affect its operational requirements.

Manufacturing defects in the Titan upgrade (OLCF-3) hardware required the rework of every blade in all 200 cabinets. This required removing portions of the system from production so that the maintenance could be performed. Performance targets were adversely affected. The maintenance was completed in December 2013. Titan seems to be operating exceptionally well now, but the risk will remain active during continued observation.

5.8.9 ID# 979—Insufficient Funding to Meet DOE Commitments (FY2013)

Risk Owner	Arthur S. Bland, OLCF Project Director		
Probability	Low		
Impact	***Cost:*** Low	***Schedule:*** Low	***Scope/Tech:*** Low
Rating	MEDIUM		
Status	**RETIRED**. Risk did not occur. Funding was adequate.		

Annual budgets are set with guidance from the ASCR office, but actual allocated funds are unknown until Congress passes funding bills. Continuing resolutions are common, and often several months will pass before actual funding is resolved. The risk is that the center may have to delay some purchases, activities, hiring, etc., or adjust lease payment schedules, resulting in high costs or schedule delays.

As the year progressed, this risk rating was reduced from Medium to Low until the risk was eventually retired. Funding was sufficient for FY2013. The OLCF maintained close contact with the federal project director and ASCR Program Office to understand the changing funding projections so that alternate plans could be made in sufficient time. Where possible, the center structured contracts to accommodate flexible payment terms.

5.8.10 ID# 1001—Leadership Computing Performance Target Is Not Achieved (CY2013)

Risk Owner	Arthur S. Bland, OLCF Project Director		
Probability	Low		
Impact	***Cost:*** Low	***Schedule:*** Low	***Scope/Tech:*** Low
Rating	LOW		
Status	**RETIRED:** The risk did not occur.		

Application readiness may be insufficient to meet the demands of the increased complexity of the new heterogeneous system architecture, resulting in too many jobs submitted that do not achieve "leadership" status. Missing performance targets will disappoint sponsors and users, which may have a lasting impact in the form of reduced system use and reduced future support for the OLCF.

Trigger: Periodic performance reports showing a trend toward possible missed goals

The OLCF Scientific Computing Group and the CAAR team will continue to strive for improvement in application readiness. The OLCF has established job queue policies with a high preference for leadership jobs. The OLCF will continue involvement with the INCITE proposal selection process such that leadership-class projects receive preference.

5.8.11 ID# 913—Long Term Support for Lustre

Risk Owner	Sudharshan S. Vazhkudai, Group Leader, Technology Integration		
Probability	Low		
Impact	***Cost:*** Low	***Schedule:*** Low	***Scope/Tech:*** Low
Rating	LOW		
Status	**RETIRED**: Risk did not occur		

The acquisition of the dedicated Lustre efforts by larger companies is relevant to the issue of long-term support for Lustre, which is critical to the success of the OLCF. The lack of long-term support could have seriously hindered our ability to field a production quality file system for the OLCF machines.

Trigger: Intelligence of lost of support or another corporate acquisition threat or action.

Intel has acquired Whamcloud, and Lustre is thriving in this setup. Further, OpenSFS is ensuring the continued progress of Lustre.

5.8.12 ID# 994—Costs for Site Prep for the Disk Acquisition May Come in Higher Than Estimated in Original Project Budget

Risk Owner	Bart A. Hammontree, Project Manager, Facilities and Operations Division		
Probability	Medium		
Impact	***Cost:*** Low	***Schedule:*** Low	***Scope/Tech:*** Low
Rating	LOW		
Status	**RETIRED**: Risk did not occur		

Delays in obtaining the requirements for the disk storage system may result in costs exceeding the original budget. Costs for site prep for the disk acquisition may come in higher than estimated in the original project budget. Cost overruns in site prep would necessitate transfer of project contingency funds or reduction in scope.

Trigger: Early intelligence on probable costs and specifications.

5.8.13 ID# 1009—Impact of New Travel Regulations on the OLCF Training Program

Risk Owner	Ashley D. Barker, Group Leader, User Assistance and Outreach		
Probability	Low		
Impact	***Cost:*** Low	***Schedule:*** Low	***Scope/Tech:*** Low
Rating	LOW		
Status	Mitigating the risk		

Due to stricter travel regulations, OLCF users who are members of DOE facilities may not be able to travel to take advantage of on-site OLCF workshops. In addition, the OLCF will have to limit on-site participation to avoid going over the newly established $100K threshold. If fewer people participate in training, there could be an impact to user satisfaction, the leadership metric, and utilization.

Trigger: Low attendance at training sessions.

The risk did not occur. The center was able to mitigate the risk through webcasting, so anyone who wanted to participate in OLCF workshops could do so free of charge. The center had the same number or more of participants at our workshops this year between on-site and webcast attendees.

5.8.14 ID# 721—Lustre Metadata Performance Continues to Impact Applications

Risk Owner	Sudharshan S. Vazhkudai, Group Leader, Technology Integration		
Probability	Medium		
Impact	***Cost:*** Low	***Schedule:*** Low	***Scope/Tech:*** Low
Rating	LOW		
Status	Mitigating the risk		

Metadata performance is critical to a wide variety of leadership applications. Its performance depends on many factors, all of which need to be optimized. Lustre performance has been stymied by not being able to scale beyond a single server and by limited performance on the server. There is a risk that single metadata server performance will not be adequate and may adversely affect both applications and interactive users. This risk has already occurred and will continue affecting performance.

Trigger: Direct observations reported by users or staff

The OLCF is working with other major Lustre stakeholders through OpenSFS to develop features to improve single metadata server performance and follow-on support of multiple metadata servers for the Lustre file system. The center has deployed Lustre 2.4, which has the Distributed Namespace (DNE) feature meant to alleviate the metadata bottleneck. The center will turn on the DNE feature in early spring of 2014, after which it will be able to determine whether DNE alleviates the metadata performance bottleneck. Multiple file systems have been deployed, reducing load on the metadata server.

6. SITE OFFICE SAFETY METRICS

CHARGE QUESTION 6: Has the facility incorporated site office safety recommendations appropriately?

OLCF RESPONSE: Yes.

6.1 SUMMARY

Oak Ridge National Laboratory (ORNL) is committed to operating under the DOE safety regulations specified in 10 C.F R. 851,[***] "Worker Safety and Health Program" as well as applicable DOE Orders and Standards. These safety requirements are incorporated into the ORNL contract, as required compliance documents. To implement these safety requirements in a consistent manner across ORNL, UT-Battelle deploys an online procedure management system referred to as the Standards-Based Management System (SBMS). Within SBMS, there are work control requirements that describe the processes to be used within ORNL operations and R&D to implement integrated safety management (ISM) functions and principles. The use of ORNL's ISM process culminates with the development and implementation of research safety summaries (RSSs), which are reviewed and approved by the ORNL Safety Services Division, line managers, and the research staff.

An RSS provides the means by which ORNL management and staff can plan and conduct research in a safe manner. It is used to control work, train participants, and provide information regarding operations and emergency services if ever needed. Under a work control review system, work plans are also written before maintenance work is allowed to proceed, to ensure that the work is conducted safely. Safety specifications are written into the service contracts and undergo a review by the authority having jurisdiction (AHJ) before new-construction and service subcontractors are allowed to begin work.

Safety assessments are conducted on RSSs, work plans, and subcontracts as well as inspections of job sites throughout each year. Lessons learned, safety snapshots, and management assessments are conducted and recorded into (ACTS. ACTS provides feedback for the completion of the ORNL ISM process. The DOE ORNL Site Office (OSO) participates in the field implementation and documentation of all of the operational safety reviews, and also partners with the ORNL Offices of Institutional Planning and Integrated Performance Management and Safety Service Division on some independent safety management system assessments.

The culture of safety at ORNL is reflected in the above processes, which seek to reduce and prevent injuries to our personnel and their potential exposure to hazards associated with the operation of the facility. The OLCF works closely with the OSO and Regina Chung, the Federal Project Director, who

[***] 10 C.F.R. 851 outlines the requirements for a worker safety/health program to ensure that DOE contractors and their workers operate a safe workplace. Additionally, 10 C.F.R. 851 establishes procedures for investigating whether a violation of a requirement of this part has occurred, for determining the nature and extent of any such violation, and for imposing an appropriate remedy.

solicited the following feedback from the OSO staff in the Operations & Oversight Division regarding OLCF's safety culture.

Mark Belvin, ES&H Specialist, *"The Site Office subject matter experts for safety and environment are part of the OLCF team. They review documentation provided by the contractor staff and perform joint walkthroughs of the project as it progresses with the DOE federal project director and contractor staff. Any issues identified are provided to the contractor for action."*

Ruth Miller, Safety and Occupational Health Specialist *"I recall that on past walkthroughs at subject project site with you we had observed work and expressed concerns over how hearing conservation, electrical safety, slip-trip-fall hazards, as well as construction safety issues while moving large equipment into small spaces, proper labeling of equipment, labeling of electrical sources, and use of appropriate personal protective equipment for the task at hand were being addressed to ensure a safe and healthful work environment for all. As I recall, each time, our questions and concerns were answered and/or addressed quickly (while we were present at the site—or shortly thereafter). I felt that significant effort was made to ensure the safety for all personnel and that compliance with 10CFR851, 29CFR1910, and 29CFR1926 standards was evident."*

David Carden, Assessment Program Manager, *"The ORNL Site Office (OSO) provides oversight for the operations of ORNL. In FY2013, the OSO Operations & Oversight Division conducted an assessment of UT-B's Contractor Assurance System, the Computing and Computational Sciences Directorate was one of the 10 organizations selected for the evaluation. Jim Hack was included in the assessment interviews. The preliminary results concluded that 'Although the main purpose for the existence of ORNL is R&D, a commendable level of service is provided for the review, oversight, and assurance of quality, safety, and environmental performance. Such service is critical to ensuring that operational performance remains safe, complaint, and effective and does not lead to issues adversely affecting ES&H, quality, or the delivery of impactful R&D products.'"*

7. CYBER SECURITY

CHARGE QUESTION 7: Has the site been certified to operate (cyber security)?

OLCF RESPONSE: Yes, the most recent OLCF authority to operate (ATO) was granted on June 17, 2013. The current ATO expires on June 16, 2014.

7.1 SUMMARY

All information technology (IT) systems operating for the federal government must have certification and accreditation (C&A) to operate. This involves the development of a policy, the approval of the policy, and an assessment of how well the organization is managing those IT resources—an assessment to determine that the policy is being put into practice.

The OLCF has the authority to operate for 1 year under the ORNL C&A package approved by DOE on June 17, 2013. The ORNL C&A package uses *Recommended Security Controls for Federal Information Systems and Organizations* [National Institute of Standards and Technology Special Publication 800-53, revision 3 (2009)] as a guideline for security controls. The OLCF is accredited at the moderate level of controls, which authorizes the facility to process sensitive, proprietary, and export-controlled data.

In the future, it is inevitable that cyber security planning will become more complex as the center continues in its mission to produce great science. As the facility moves forward, the OLCF is very proactive, viewing its cyber security plans as dynamic documentation and responding to and making modifications as the needs of the facility change to provide an appropriately secure environment.

Summary of the Proposed Metric Values

HIGH PERFORMANCE COMPUTING FACILITY
2013 OPERATIONAL ASSESSMENT
OAK RIDGE LEADERSHIP COMPUTING FACILITY

February 2014

8. ACTUAL AND PROPOSED METRIC VALUES

CHARGE QUESTION 8: Are the performance metrics used for the review year and proposed for future years sufficient and reasonable for assessing operational performance?

OLCF RESPONSE: Yes. The OLCF works closely with the DOE program manager to develop and update metrics and to target values that reflect the expectations of the stakeholders in delivering a leadership-class HPC resource.

8.1 SUMMARY

Table 8.1 provides a summary of the metrics and actuals for 2013 and proposed metrics and targets for 2014 and 2015.

Table 8.1. OLCF Metrics and Actuals for 2013, and Proposed Metrics and Targets for 2014 and 2015

2013 Metric and Target	2013 Actual	2014 Metric	2014 Target	2015 Target	Reporting Period
Are the processes for supporting the customers, resolving problems, and Outreach effective?					
Customer Metric 1: Customer Satisfaction					
Overall score on the OLCF user survey. Target: Results will be satisfactory (3.5/5.0) based on a statistically meaningful sample.	The OLCF exceeded the metric target: 4.4/5.0.	Overall score on the OLCF user survey.	Results will be satisfactory (3.5/5.0) based on a statistically meaningful sample.	Results will be satisfactory (3.5/5.0) based on a statistically meaningful sample.	Annual
Improvement on results that scored below satisfactory in the previous period. Target: Results will show improvement in at least ½ of questions that scored below satisfactory (3.5) in the previous period.	The OLCF exceeded the metric target: No question scored below satisfactory (3.5/5.0) on the 2013 survey.	Improvement on results that scored below satisfactory in the previous period.	Results will show improvement in at least one-half of the questions that scored below satisfactory (3.5) in the previous period.	Results will show improvement in at least one-half of the questions that scored below satisfactory (3.5) in the previous period.	Annual
Customer Metric 2: Problem Resolution					
OLCF survey results related to problem resolution. Target: Results will be satisfactory (3.5/5.0) based on a statistically meaningful sample.	The OLCF exceeded the metric target: 4.4/5.0.	OLCF survey results related to problem resolution.	Results will be satisfactory (3.5/5.0) based on a statistically meaningful sample.	Results will be satisfactory (3.5/5.0) based on a statistically meaningful sample.	Annual
OLCF user problem resolution time period. Target: 80% of OLCF user problems will be addressed within three business days, by either resolving the problem or informing the user how the problem will be resolved.	The OLCF exceeded the metric target: 92.3%.	OLCF user problem resolution time period.	Eighty percent of OLCF user problems will be addressed within three business days, by either resolving the problem or informing the user how the problem will be resolved.	Eighty percent of OLCF user problems will be addressed within three business days, by either resolving the problem or informing the user how the problem will be resolved.	Monthly

Table 8.1. (continued)

2013 Metric and Target	2013 Actual	2014 Metric	2014 Target	2015 Target	Reporting Period
Customer Metric 3: User Support					
Average of user support ratings. Target: Results will be satisfactory (3.5/5.0) based on a statistically meaningful sample.	The OLCF exceeded the metric target: 4.4/5.0.	Average of user support ratings.	Results will be satisfactory (3.5/5.0) based on a statistically meaningful sample.	Results will be satisfactory (3.5/5.0) based on a statistically meaningful sample.	Annual
Is the facility maximizing the use of its HPC systems and other resources consistent with its mission?					
Business Metric 1: System Availability (for a period of one year following a major system upgrade, the targeted scheduled availability is 85% and overall availability is 80%)[a]					
Scheduled Availability. Target: Titan: 85% (lower in FY12 due to the compute system upgrades); HPSS: 95%; External File Systems: 95%.	The OLCF exceeded the metric target. Titan: 98.70%; HPSS: 99.99%; Widow0: 100%; Widow1: 99.87%; Widow2: 99.91%; Widow3: 99.87%.	Scheduled availability.	90%	90%	Monthly
Overall Availability. Target: Titan: 80%; HPSS 90%; External File Systems 90%.	The OLCF exceeded the metric target: Titan: 93.82%; HPSS: 97.6%; Widow0: 99.11%; Widow1: 97.35%; Widow2: 97.52%; Widow3: 98.09%.	Overall availability.	Titan: 85%; HPSS 90%; External File Systems: existing, 90%	Titan: 85%; HPSS 90%; External File Systems: existing, 90%	Monthly
Business Metric 2: Capability Usage					
OLCF will report on capability usage. Target: In the first year of production, at least 30% of the consumed node-hours will be from jobs requesting 20% or more of the available compute nodes.	The capability usage was 59.38%. The OLCF exceeded the metric target.	OLCF will report on capability usage.	In subsequent years, at least 35% of the consumed node-hours will be from jobs requesting 20% or more of the available compute nodes.	In subsequent years, at least 35% of the consumed node-hours will be from jobs requesting 20% or more of the available compute nodes.	Monthly
N/A[b]	N/A[b]	OLCF will report GPU usage (reference only, no target).	N/A[b]	N/A[b]	Monthly

[a] The Cray XK7, Titan, went into production on May 31, 2013. The external file system, Atlas, went into production on October 3, 2013.
[b]Not applicable.

APPENDIX A. SURVEY RESPONSE RATE

Table A-1. Day-by-Day Survey Response Rate, 2013[a]

Survey Time Line	Date	Day	Number of Respondents	Percent of responses (N = 1,232)	Description of Reminder
Day 1	2-Oct	Wed	73	5.93%	Initial e-mail invitation sent by ORISE evaluator, Dr. Erin Burr.
Day 2	3-Oct	Thurs	19	1.54%	--
Day 3	4-Oct	Fri	13	1.06%	Survey link sent to users in Friday Announcements.
Day 4	5-Oct	Sat	1	0.08%	--
Day 5	6-Oct	Sun	1	0.08%	--
Day 6	7-Oct	Mon	10	0.81%	--
Day 7	8-Oct	Tues	2	0.16%	--
Day 8	9-Oct	Wed	2	0.16%	--
Day 10	11-Oct	Fri	8	0.65%	Survey link sent to users in Friday Announcements.
Day 11	12-Oct	Sat	2	0.16%	--
Day 12	13-Oct	Sun	2	0.16%	--
Day 13	14-Oct	Mon	3	0.24%	--
Day 14	15-Oct	Tues	1	0.08%	--
Day 15	16-Oct	Wed	1	0.08%	--
Day 16	17-Oct	Thurs	2	0.16%	--
Day 17	18-Oct	Fri	2	0.16%	Survey link sent to users in Friday Announcements.
Day 18	19-Oct	Sat	1	0.08%	--
Day 20	21-Oct	Mon	44	3.57%	E-mail reminder sent from Jack Wells, Dir. of Science, NCCS
Day 21	22-Oct	Tues	15	1.22%	--
Day 22	23-Oct	Wed	3	0.24%	--
Day 23	24-Oct	Thurs	4	0.32%	--
Day 24	25-Oct	Fri	7	0.57%	Survey link sent to users in Friday Announcements.
Day 27	28-Oct	Mon	3	0.24%	--
Day 28	29-Oct	Tues	1	0.08%	--
Day 29	30-Oct	Wed	1	0.08%	--
Day 30	31-Oct	Thurs	5	0.41%	Jack Wells made personal contacts with several non-responders.
Day 31	1-Nov	Fri	2	0.16%	Survey link sent to users in Friday Announcements.
Day 34	4-Nov	Mon	1	0.08%	--
Day 36	6-Nov	Wed	1	0.08%	--
Day 37	7-Nov	Thurs	13	1.06%	E-mail reminder sent by Buddy Bland, OLCF Project Director.

Table A-1. (continued)

Survey Time Line	Date	Day	Number of Respondents	Percent of responses (N = 1,232)	Description of Reminder
Day 38	8-Nov	Fri	6	0.49%	Survey link sent to users in Friday Announcements.
Day 40	10-Nov	Sun	1	0.08%	--
Day 41	11-Nov	Mon	1	0.08%	--
Day 42	12-Nov	Tues	1	0.08%	--
Day 44	14-Nov	Thurs	2	0.16%	--
Day 45	15-Nov	Fri	2	0.16%	Survey link sent to users in Friday Announcements.
Day 46	16-Nov	Sat	1	0.08%	--
Day 51	21-Nov	Thurs	1	0.08%	--
Day 52	22-Nov	Fri	1	0.08%	Survey link sent to users in Friday Announcements.
Day 53	23-Nov	Sat	2	0.16%	--
Day 55	25-Nov	Mon	16	1.30%	--
Day 56	26-Nov	Tues	7	0.57%	--
Day 57	27-Nov	Wed	3	0.24%	--
Day 59	29-Nov	Fri	1	0.08%	Survey link sent to users in Friday Announcements.
Day 61	1-Dec	Sat	2	0.16%	--
Day 62	2-Dec	Sun	12	0.97%	--
Day 63	3-Dec	Mon	49	3.98%	Final reminder email sent by Jack Wells, Dir. of Science, NCCS
Day 64	4-Dec	Tues	7	0.57%	--
Day 65	5-Dec	Wed	1	0.08%	--
Day 66	6-Dec	Thurs	8	0.65%	--
		Total	**367**	**29.79%**	

[a]Days during which no OLCF users responded to the survey are not included in this time line.

APPENDIX B. OUTREACH PRODUCTS

Table B-1. Outreach Product Listing, 2013

Date	Type of Product	Title
1/2	Misc. Publication	What is a Core-Hour on Titan? Factsheet
1/2	Misc. Publication	Oak Ridge Leadership Computing Facility Factsheet
1/2	Misc. Publication	National User Facility Organization Poster
1/14	Highlight	Sea-Level Rise Will Continue Even with Aggressive Emission Mitigation
1/14	Quad Chart	Sea-Level Rise Inevitable
1/14	Highlight	ORNL's EVEREST Upgraded
1/14	PPT Slide	Oak Ridge National Laboratory's EVEREST Upgraded
1/14	Highlight	ORNL Continues Strong Leadership Tradition at 2012 Supercomputing Conference
1/14	PPT Slide	ORNL Continues Strong Leadership Tradition at 2012 Supercomputing Conference
2/18	Websites	OLCF Migration (Migration of OLCF website to new servers)
2/18	Highlight	Lessons From the Past
2/18	Quad Chart	Climate Science: Ocean circulation drove warming and deglaciation of Southern Hemisphere
2/18	Highlight	Extermination at Scale
2/18	PPT Slide	Extermination at Scale
2/18	Highlight	User Conference Brings Titan Training to the West Coast
2/18	PPT Slide	User Conference Brings Titan Training to the West Coast
3/18	Highlight	Supernovas as Nuclear Pasta Factories
3/18	Quad Chart	Simulations Map Nuclear Pasta Phase in a Core-Collapse Supernova
3/18	Highlight	Titan Users Now Have Access to GPUs
3/18	PPT Slide	Titan Users Now Have Access to GPUs
3/18	Highlight	Workshop Prepares Users to Run on Titan
3/18	PPT Slide	Workshop Prepares Users to Run on Titan
4/11	Highlight	ORNL High-Performance Computing Team Leader Honored for Career Achievements
4/15	Highlight	Call for Research Proposals at U.S. Leadership Computing Facility for Advances in Science and Engineering
4/29	Highlight	OLCF Heads West
4/29	PPT Slide	OLCF Heads West
4/29	Highlight	Titan Shows Life Science Advancements at Bio-IT Conference
4/29	PPT Slide	Titan Shows Life Science Advancements at Bio-IT Conference
5/20	Highlight	Jaguar Guides Demonstration of Novel Quantum State
5/20	Quad Chart	Simulations Lead to Experimental Verification of Bose Glass
5/20	Highlight	OLCF Announces New Head of Scientific Computing
5/20	PPT Slide	OLCF Announces New Head of Scientific Computing
5/20	Highlight	OLCF Industry Partnerships Highlighted in Fortune Article
5/20	PPT Slide	OLCF Resources Highlighted in Fortune Article
5/20	Highlight	Lattice QCD Community Visits ORNL to Discuss Challenges in the Field
5/20	PPT Slide	Lattice QCD Community Visits ORNL to Discuss Challenges in the Field
5/20	Highlight	Titan User Recognized by the American Chemical Society
5/20	PPT Slide	Titan User Recognized by the American Chemical Society

Table B-1. (continued)

Date	Type of Product	Title
6/10	Highlight	Hold It Right There
6/10	Quad Chart	CFD Modeling and Simulation for CASL using Hydra-TH
6/10	Highlight	Titan Completes Acceptance Testing
6/10	PPT Slide	Titan Completes Acceptance Testing
6/10	Highlight	Moniz Visits ORNL
6/10	PPT Slide	Moniz Visits ORNL
6/10	Highlight	ORNL Provides Leadership at 2013 Lustre User Group Conference
6/10	PPT Slide	ORNL Provides Leadership at 2013 Lustre User Group Conference
6/10	Highlight	OLCF Representatives Make Impact at Cray User Group Meeting in California
6/10	PPT Slide	OLCF Makes Impact at Cray User Group Meeting in California
6/10	Websites	StatusCast dashboard to display OLCF resource statuses
6/10	Websites	Transition of NCCS website to new ORNL content management system
7/8	Highlight	Vampir Takes a Bite Out of Inefficiency as Codes Run on Bigger Supercomputers
7/8	PPT Slide	Enabling Titan Sized Code Development
7/8	Highlight	OLCF User Earns Early Honor
7/8	PPT Slide	OLCF User Earns Early Honor
7/8	Highlight	ADIOS Wins Big
7/8	PPT Slide	ADIOS Wins Big
7/8	Highlight	ORNL Representatives Travel to Germany
7/8	PPT Slide	ORNL Representatives Travel to Germany
7/8	Highlight	ORNL Paper Brings Home International Award
7/8	PPT Slide	ORNL Paper Brings Home International Award
7/8	Poster	Oak Ridge Leadership Computing Facility Recruiting Poster
7/25	Highlight	Early Molecular Dynamics Research Blazes through Titan's New GPUs
8/12	Highlight	ORNL's Jaguar Gets Under the Hood
8/12	Quad Chart	ORNL's Jaguar Gets Under the Hood
8/12	Quad Chart	All SC13 Gordon Bell Finalists Use DOE Supercomputers
8/12	Quad Chart	Four of Six Gordon Bell Finalists Use Titan
8/12	Highlight	Pushing the Envelope
8/12	Highlight	Supersizing Spider
8/12	PPT Slide	Supersizing Spider
8/12	Highlight	OpenACC Headlines OLCF Summer Workshops
8/12	PPT Slide	OpenACC Headlines OLCF Summer Workshops
8/12	Highlight	High Honors For OLCF Users
8/12	PPT Slide	High Honors For OLCF Users
8/12	Highlight	OLCF in the News
8/12	PPT Slide	OLCF in the News
8/20	Highlight	IEEE Computing in Science & Engineering Special Issue on Leadership Computing Announced
9/3	Websites	ACCEL Webpage
9/3	Highlight	Titan Sheds Light on Unknowns in Organic Photovoltaic Research
9/3	Quad Chart	Towards Rational Design of Efficient Organic Photovoltaic Materials
9/3	Highlight	Users Defeat 'Data Deluge' at OLCF's First Large Data Sets Workshop
9/3	PPT Slide	OLCF Hosts First Large Data Sets Workshop
9/3	Highlight	OLCF Continues Education Leadership with Summer Students

Table B-1. (continued)

Date	Type of Product	Title
9/3	PPT Slide	OLCF Continues Education Leadership with Summer Students
9/3	Highlight	OLCF Web Developer Receives Professional Honor
9/3	PPT Slide	OLCF Web Developer Receives Professional Honor
9/23	Highlight	Titan's New Build Attracts Magnetic Systems Research Impossible Until Now
9/23	Quad Chart	Magnetic Materials
9/23	Highlight	ForWarn Researchers Get EVEREST-Sized Look at Woodland Disturbances
9/23	Quad Chart	ForWarn Researchers Get EVEREST-Sized Look at Woodland Disturbances
9/23	Highlight	Oak Ridge Leadership Computing Facility User Update: SmartTruck Systems
9/23	Quad Chart	Computational Fluid Dynamics "Smart Truck Optimization"
9/23	Highlight	ADIOS Code Sprint: A Race for New Technologies
9/23	PPT Slide	ADIOS Code Sprint: A Race for New Technologies
9/23	Poster	Electronic Posters: ACCEL and OLCF Electronic Posters for SC13
10/14	Publication	Acceleration Competitiveness through Computational Excellence (ACCEL) brochure
10/21	Highlight	Titan Propels GE Wind Turbine Research into New Territory
10/21	Quad Chart	Non-Icing Surfaces for Wind Turbines
10/21	Highlight	OLCF Staff to Win Big at Awards Night
10/21	PPT Slide	OLCF Staff to Win Big at Awards Night
10/21	Highlight	Titan Gets Computational Company
10/21	PPT Slide	Titan Gets Computational Company
10/21	Highlight	The Future of Computing and Data Integration
10/21	PPT Slide	The Future of Computing and Data Integration
10/21	Websites	Maintenance and upgrade of NCRC website. Migration of NCRC website to new servers.
11/4	Highlight	Peering into Cells One GPU at a Time
11/4	Quad Chart	Peering into Cells One GPU at a Time
11/4	Highlight	Code for Largest Cosmological Simulations Ever on GPUs Is Gordon Bell Finalist
11/4	Quad Chart	Exploring the Nature of the Lightest Massive Particles in the Universe
11/4	Highlight	Protecting Big Data
11/4	PPT Slide	Protecting Big Data
11/4	Highlight	High-Risk, High-Reward Simulations
11/4	Quad Chart	High-Risk, High-Reward Simulations
11/12	Highlight	Simulations of Plasma Turbulence Model the Inner Workings of Cosmic Phenomenon
11/12	Quad Chart	Turbulent Plasma Simulations Using PICONGPU
11/12	Highlight	Superconductor Simulation Tops 15 Petaflops On Titan
11/12	Quad Chart	Superconductor Simulation Tops 15 Petaflops On Titan
11/14	Highlight	OLCF Lends Expertise for Introducing GPU Accelerator Programming to Popular Linux GCC Compiler
11/18	Websites	SC13 Website
11/18	Websites	OLCF HPC Blog
11/18	Press Release	INCITE grants awarded to 59 computational research projects
11/19	Highlight	Oak Ridge Leadership Computing Facility receives three HPCwire awards
11/20	Highlight	Four OLCF Partners Win Major HPC Award
11/25	Poster	INCITE 2014 Projects Poster
12/2	Publication	Oak Ridge Leadership Computing Facility Annual Report 2012-13
12/2	Highlight	Titan Simulates Earthquake Physics Necessary for Safer Building Design

Table B-1. (continued)

Date	Type of Product	Title
12/2	Quad Chart	Physics-based Probabilistic Seismic Hazard Analysis
12/2	Highlight	Simulation Shuffles Protons and Electrons
12/2	Quad Chart	Coupled Electronic and Nuclear Dynamics in Solar Photocatalytic Water Splitting
12/2	Highlight	OLCF Wins Big at Conference
12/2	PPT Slide	OLCF Wins Big at Conference
12/2	Highlight	ORNL Intern Wins Best Abstract Award at Student Poster Session
12/2	PPT Slide	ORNL Intern Wins Best Abstract Award at Student Poster Session
12/16	Highlight	Researchers recruit Titan to study key molecular switch that controls cell behavior
12/16	Quad Chart	Discovery and characterization of a trans-membrane molecular switch
12/16	Highlight	Boosting Bioenergy and Overcoming Recalcitrance
12/16	Quad Chart	Boosting Bioenergy and Overcoming Recalcitrance
12/16	Highlight	The Need for Speed
12/16	PPT Slide	The Need for Speed
12/16	Highlight	Out with the Old, In with the New
12/16	PPT Slide	Out with the Old, In with the New
12/16	Poster	ALCC Projects Poster 2013-14

APPENDIX C. TRAINING

Table C-1. Training List, 2013

Event Type	Description	Date	Participants
Workshop/Training	Titan Training West Coast	Jan 29-31	60
Workshop/Training	OOP (Object-Oriented Programming) Workshop	Feb 5-7	68
Seminar Series	T.P. Straatsma, Pacific Northwest National Laboratory, “Eliminating Explicit Synchronization from Molecular Simulations to Improve Scalability”	Feb 7	Unknown
Workshop/Training	Titan Training East Coast	Feb 19-21	116
Seminar Series	Filippo Spiga, University of Cambridge, “Quantum ESPRESSO”	Mar 15	Unknown
Seminar Series	Joost VandeVondele, ETH Zurich, “Density Functional Theory (DFT) Based Simulation”	Apr 5	Unknown
Workshop/Training	LQCD Workshop	Apr 29–May 3	25
User Con Call	May User Conference Call	May 7	3
Workshop/Training	INCITE Proposal Writing Webinar	May 14	49
Seminar Series	Thomas Wuest, Swiss Federal Research Institute WSL, “Unknotting Challenging Questions in Protein Physics using Wang-Landau Sampling”	May 23	Unknown
User Con Call	June User Conference Call	June 4	5
Workshop/Training	Programming with Big Data in R: pbdR	June 17	22
Workshop/Training	Crash Course in Supercomputing	Jun 18–19	45
Seminar Series	Karol Kowalski, Pacific Northwest National Laboratory, “The Coupled Cluster Formalism Across Spatial, Energy, and Time (to solution) Scales”	June 27	Unknown
User Con Call	July User Conference Call	July 2	3
Workshop/Training	OpenACC Tutorials	Jul 15–18	86
Workshop/Training	GPU Programming for Molecular Modeling Workshop	Aug 3	12
Workshop/Training	Processing and Analysis of Very Large Data Sets	Aug –8	80
Seminar Series	Michael S. Pindzola, Auburn University, “Atomic and Molecular Collisions using a Time-Dependent Close-Coupling Method”	Aug 16	Unknown
Workshop/Training	ADIOS Code Spring Workshop	Aug 19–20	19
Internal Staff Training	NVIDIA Train the Trainer Workshop	Aug 20–21	24
User Con Call	September User Conference Call	Sep 3	29
Seminar Series	Stan Tomov, The University of Tennessee, “High-Performance Linear Algebra with Intel Xeon Phi (MIC) Coprocessors”	Sep 9	Unknown
Workshop/Training	OpenACC Face-to-Face Meeting	Sep 24-2–6	30
User Con Call	October User Conference Call	Oct 1	33
Workshop/Training	XC30 "Eos" Training	Oct 17	28
Seminar Series	Patrick Charbonneau, Duke University “High-Dimensional Surprises Near the Glass and the Jamming Transitions”	Oct 28	Unknown
User Con Call	November User Conference Call	Nov 5	29
Seminar Series	Rebecca Hartman-Baker, iVEC “Enigmas in the Outback: Computational Science at Unprecedented Scales”	Dec 3	Unknown
Seminar Series	Andreas Wingen, Oak Ridge National Laboratory, “Regularization of Soft-X-ray Imaging in the DIII-D Tokamak”	Dec 12	Unknown

APPENDIX D. 2013 DIRECTOR'S DISCRETIONARY ALLOCATIONS

Table D-1. 2013 Director's Discretionary Allocations

PI	Affiliation	2013 Allocation	2013 Usage	Project Name
Michael Bussmann	Helmholtz-Zentrum Dresden-Rossendorf	20,000,000	19,260,895	Laser-Wakefield Simulations Using PICONGPU
Rainald Lohner	George Mason University	1,470,000	69,414	Highly Detailed Simulations of Blasts on Offshore Platforms
Dominic von Terzi	GE Global Research	5,880,000	1,964,682	LES for Wind Turbine Interactions
Dana Hammond	NASA-LARC	32,687	436,344	Scaling of NASA CFD Application for Aeronautics
Vittorio Michelassi	GE Global Research	4,410,000	889,387	HIPSTAR-G
Guglielmo Scovazzi	Duke University	2,000,000	0	Athena-VMS: A new transient dynamic framework for complex geometry computations
Dana Hammond	NASA-LARC	1,000,000	478,174	Scaling of FUN3D for 2014 INCITE Proposal
Allan D Grosvenor	Ramgen Power Systems	1,000,000	0	Vizualization of Tip Injection Phenomena in the near Stall Regime of Transonic Fan Stage
Allan D Grosvenor	Ramgen Power Systems	2,000,000	2,444,918	Compressible Flow Turbomachinery Optimization: Numerical Tools Advancement
Bronson Messer	ORNL	11,760,000	2,952,171	Explosive Nucleosynthesis and Deflagration to Detonation in Type Ia Supernovae
Patrick Fragile	ORAU	64,795	0	Radiation Transport in Numerical Simulations of Black-Hole Accretion Disks
Michael Warren	LANL	2,000,000	11,157	The Dark Sky Simulations
Tiziana Di Matteo	Carnegie Mellon University	2,940,000	276,275	Petascale Cosmology with P-Gadget
Simon Portegies Zwart	Leiden University	2,000,000	1,964,439	The Fine Structure of the Milky Way Galaxy
Michael Clark	NVidia	1,000,000	887,485	Petascale Cross Correlation
Zhihui Du	Tsinghua University	1,000,000	1,640,801	Gravitational Wave Source Modeling based on Binary Black Hole Simulation
Adam Burrows	Princeton University	1,000,000	256,326	3D Radiation Hydrodynamics Simulations of Core Collapse Supernovae
Alexander Sandor Szalay	Johns Hopkins University	60,000	0	Demonstrations of Data-Scope at 100 Gbps Across a National Data-Intensive Computational Science Test Bed at SC13

Table D-1. 2013 Director's Discretionary Allocations (continued)

PI	Affiliation	2013 Allocation	2013 Usage	Project Name
Moetasim Ashfaq	ORNL	10,000,000	12,857,641	A hierarchical regional modeling framework for decadal-scale hydro-climatic predictions and impact assessments
John Michalakes	NREL	500,000	0	Simulator for Offshore Wind Plant Applications (SOWFA)
Jason Hill	University of Minnesota	1,176,000	0	Air Quality Impacts of Conventional and Alternative Energy for Transportation
Balaji Jayaraman	Pennsylvania State University	1,000,000	284	Towards High-fidelity Petascale Computations of Atmospheric Turbulence-driven Wind Turbine Aerodynamics and Wakes
Jimy Dudhia	NCAR	200,000	151,341	MPAS DYNAMO
Xin-Zhong Liang	University of Maryland College Park	100,000	21,284	Benchmarking for CWRF Climate Prediction
Richard Casey	Colorado State University	250,000	98,259	Large-Scale Metagenomic and Bioinformatic Data Analysis of Semiconductor-based Next Generation DNA Sequencing
Jerome Baudry	UT-Knoxville	12,350,000	3,143,617	Massive ensemble docking for drug toxicity prediction
Chongle Pan	ORNL	4,440,000	3,884,787	Large-scale metagenomics analysis for biosurveillance and environmental microbiology
Gustavo Seabra	Universide de Federal de Pernambuco	561,467	118,491	Elucidation of the Molecular Mechanism of Enzymatic Reactions by Molecular Dynamics and Hybrid Quantum Mechanical and Molecular Mechanics Simulations
Giuseppe Milano	Universita degli Studi di Salerno	50,000	0	GPU Accelerated Hybrid Particle Field Molecular Dynamics Simulations
Miguel Fuentes-Cabrera	ORNL	3,000,000	1,070,717	Theoretical study on the Molecular Transport across Bacterial Micro-compartments
Xiaolin Cheng	ORNL	2,000,000	421,533	Computational Study of Cellulose Synthase via Enhanced Sampling in High Performance Computing
Rommie Amaro	University of California San Diego	2,000,000	775,067	Simulation of Large-Scale Biomolecular Systems
Karen Lee Russ	University of Wisconsin	2,000,000	4,120,216	Disruption of lamellar lipid systems induced by small molecule permeants
Nikolay Dokholyan	University of North Carolina Chapel Hill	500,000	16,589	Characterization of structure and dynamics of clinically relevant membrane proteins by means of molecular dynamics simulations

Table D-1. 2013 Director's Discretionary Allocations (continued)

PI	Affiliation	2013 Allocation	2013 Usage	Project Name
Pratul K Agarwal	ORNL	1,000,000	967,347	Characterizing the Conformational Sub-states for Developing Hyper-catalytic Enzymes
Chris Mundy	PNNL	11,760,000	9,709,747	Control of Complex Transformations with Advanced Molecular Simulation
Greg Voth	University of Utah	1,500,000	572,554	Energy Storage and Conversion Materials
Theresa Windus	Iowa State University	500,000	371,519	Critical Materials Institute: Separations Science
Erik Deumens	University of Florida	1,290,567	1,085,928	EOM-CC calculations on diamond nano crystals
James Joseph Hack	ORNL	23,520,000	10,249,015	Ultra High Resolution Global Climate Simulation to Explore and Quantify Predictive Skill for Climate Means, Variability and Extremes
Katherine Evans	ORNL	1,000,000	475,847	A Scalable, Efficient, and Accurate Community Ice Sheet Model (SEACISM)
Colin Jones	Swedish Meteorological and Hydrological InstituteFolkborgsvagen	1,470,000	0	HIRES-CORDEX
Thomas Henderson	NOAA-GFDL	12,272	55,432	GPU Computing for Numerical Weather Prediction
Robert Cook	ORNL	29,400	0	Modeling and Synthesis Thematic Data Center (MAST-DC)
Salil Mahajan	ORNL	2,000,000	28,490	Impact of Aerosols and Air-sea Interactions on CESM Biases in the Western Pacific Warm Pool Region
Dali Wang	ORNL	200,000	726	Parallel Geospatial Data Management for Multiscale Environmental Data Analysis on GPUs
Ramanan Sankaran	ORNL	735,000	55	Simulating combustion in automotive engines with real fuel chemistry
Tang-Wei Kuo	General Motors	1,470,000	485,666	Multi-hole injector optimization for spark-ignited direct-injection gasoline engines
Brad VanDerWege	Ford Motor Company	1,470,000	274	Cycle-to-Cycle Combustion Variation Modeling
Sreekanth Pannala	ORNL	735,000	921,964	Computational Infrastructure for parallel simulations of Cycle-to-Cycle variations of in-cylinder combustion
Sreekanth Pannala	ORNL	735,000	897,780	Parallel computational infrastructure for optimizing Multi-hole injector for spark-ignited direct-injection gasoline engines

Table D-1. 2013 Director's Discretionary Allocations (continued)

PI	Affiliation	2013 Allocation	2013 Usage	Project Name
Yi Wang	ORNL	1,470,000	53,961	CFD Modeling of Industrial Scale Fire Growth and Suppression
John Bell	LBNL	2,000,000	130	Analysis of high-fidelity simulations of premixed turbulent flames
Suresh Menon	Georgia Institute of Technology	10,000	8	Simulations of Detonation to Deflagration Transition in Two-Phase Reactive Mixture and Supercritical Combustion in High Pressure Shear Co-axial Injector
Roy Primus	GE Global Research	500,000	22	Application of High Performance Computing for Simulating Cycle to Cycle Variation in Dual Fuel Combustion Engines
Vaidyanathan Sankaran	United Technologies Research Center	1,500,000	1,458,864	Towards Combustor Simulation Using Large Eddy Simulation and Graphical Processing Units
Kritjan Haule	Rutgers University	1,000,000	1,518,145	Calculation of Strongly Correlated Systems Using DMFT(CTQMC+WIEN2K) Method
Nicola Varini	Curtin University	2,400,000	2,099,489	EXX-PETA
Jens Glaser	University of Minnesota	147,000	152,527	Optimization of a general-purpose molecular dynamics code running on multiple GPUs
Mark Oxley	Vanderbilt University	5,000,000	4,169,039	Simulation of atomic-resolution electron energy loss spectra on the meso-scale
Xiaoguang Zhang	ORNL	1,700,000	10,421	A Comprehensive Theoretical/Numerical Tool for Electron Transport in Mesoscale-Heterostructures
Stephen Poole	ORNL	0	0	FASTOS Community Allocation
Terry Jones	ORNL	3,000,000	0	HPC Colony II
Joshua New	ORNL	500,000	327,208	Autotune E+ Buildings
Kalyan Perumalla	ORNL	2,940,000	0	ReveR-SES: Reversible Software Execution Systems for Ultra-scale Computing
Barbara Chapman	University of Houston	219,449	102	A similarity-based analysis tool for pattern derivation and large-scale program restructuring
Patrick Joseph Burns	Colorado State University	0	3	Grad 511
George Biros	University of Texas Austin	7,350,000	3,572,073	Fast N-body algorithms in high-dimensions
Olaf Schenk	Universita della Svizzera italina	500,000	0	Large-Scale Seismic Imaging on HPC Architectures: Applications, Algorithms and Software

Table D-1. 2013 Director's Discretionary Allocations (continued)

PI	Affiliation	2013 Allocation	2013 Usage	Project Name
David Pugmire	ORNL	8,500,000	8,606,165	SDAV
Richard Mills	ORNL	2,940,000	178,819	Hierarchical Krylov Methods for Ultrascale Computers
Martin Burtscher	Texas State University San Marcos	73,500	9,564	GPU Application performance and data analytics
Yuji Shinano	Zuse Institute Berlin	1,470,000	1,491,509	ParaSCIP
Jim Tallman	GE Global Research	2,940,000	409,197	Tacoma Scalability for INCITE-sized problems
Andreas Schaefer	Friedrich-Alexander-Universitaet Erlangen-Nuernberg (FAU)	1,470,000	63,310	LibGeoDecomp
Rajiv Sampath	GE Global Research	0	0	GE Global Research
Adam Simpson	ORNL	100,000	1,822	Supercomputing in the Classroom
Rebecca Hartman Baker	ORNL	1,000,000	0	Efficacy of GPGPU-Accelerated System for Pawsey Centre Phase II
Sergey Panitkin	BNL	500,000	523	Next Generation Workload Management System
Fernanda Schafer Foertter	ORNL	5,000,000	163,652	Developing Scalable Heterogeneous Computing Training Code Examples
Terry Jones	ORNL	3,000,000	0	HPC Colony
Benson Muite	King Abudallah University of Science and Technology	1,000,000	896,805	Numerical investigations of semilinear partial differential equations
Oscar Hernandez	ORNL	400,000	700,164	SCALPERF
Bradley Settlemyer	ORNL	1,000,000	175,452	Towards a Resilient and Scalable Infrastructure for Big Data
Bronson Messer	ORNL	1,000,000	29	CORAL Benchmarking
Clifton Woolley	NVIDIA	1,000,000	798,329	MiniApps
Judith C Hill	ORNL	5,000,000	502,398	Computational Partnerships
Jason Micah Cope	DataDirect Networks	200,000	21,031	Assessing the Scalability of DataDirect Networks Iron Monkey Burst Buffer on Titan
Judith Hill	ORNL	100,000	3,557	Computational Science Graduate Fellowship Program
Zhihong Lin	University of California Irvine	8,820,000	3,309,764	Porting and scaling of GTC code on GPU-based architecture

Table D-1. 2013 Director's Discretionary Allocations (continued)

PI	Affiliation	2013 Allocation	2013 Usage	Project Name
Steven Shannon	North Carolina State University	1,470,000	0	Particle-In-Cell Simulation of Radio Frequency Field Structure Near Plasma Facing Antenna Components
William Tang	Princeton University	7,350,000	7,835,991	GPU-CPU Global PIC
Thomas Jenkins	Tech-X	6,000,000	2,039,589	Extended magnetohydrodynamic simulations of toroidal fusion plasmas
Jeff Candy	General Atomics	2,000,000	9	GYRO simulation of electromagnetic turbulence in tokamaks
David Green	ORNL	2,000,000	173,981	AORSA-VORPAL Coupling Development for Radio-Frequency Heating of Fusion Plasmas
Jorge Pita	Aramco	3,500,000	3,404,697	GPU-Accelerated Large-Scale Basin and Reservoir Simulation
James McClure	Virginia Tech	500,000	324,929	Accelerating Multiphase Flow Simulations in Porous Media
Dag Lohmann	KatRiskm LLC	5,000,000	593,185	Worldwide Flood Map
Sampath Kumar Gajawada	ORNL	57,600	0	TGS Titan Benchmark Runs
Balint Joo	JLab	1,176,000	2,517,664	Porting Lattice QCD Codes to Titan
Rene Bellwied	University of Houston	1,000,000	0	Thermodynamic aspects of the QCD crossover region from the lattice
Keh-Fei Liu	University of Kentucky	2,000,000	1,807,310	Proton Spin Components from Lattice QCD
Jacques Corbeil	KatRiskm LLC	250,000	240,839	Next Generation De Novo Assembler
Rong Tian	Institute of Computing Technology, Chinese Academia of Sciences	1,470,000	1,547,902	Petascale simulation of fracture process
Ashok Srinivasan	Florida State University	439,459	11,846	Accelerating Quantum Monte Carlo on Massively Parallel Computing Platforms
Predrag Krstic	ORNL	1,470,000	0	Science of the Plasma-Material Interface at Extreme Conditions
Jacek Jakowski	UT-Knoxville	2,940,000	2,954,519	Electronic structure calculation methods on accelerators
Srdjan Simunovic	ORNL	250,000	6,700	Validating Predictive Modeling of Carbon Fiber Composites In Automotive Crash Applications
James Lewis	West Virginia University	1,000,000	550,488	High-Throughput Design of Delafossite Oxide Materials for Photovoltaics
Leonid Zhigilei	University of Virginia	8,820,000	4,856,490	Atomistic simulations of laser interactions with metals

Table D-1. 2013 Director's Discretionary Allocations (continued)

PI	Affiliation	2013 Allocation	2013 Usage	Project Name
Galen Shipman	ORNL	8,000,000	279,125	Accelerating Materials Modeling with Leadership Computing
Xiaoye Li	LBNL	300,000	381,816	Next Generation Computing for X-ray Science
Bala Radhakrishnan	ORNL	3,000,000	1,717,570	Multi Objective Optimization of Microstructures
Jeongnim Kim	ORNL	3,000,000	18,526	QMC Glue
Kan-Ju Lin	Massachusetts Institute of Technology	10,000	0	Radiation tolerance and mechanical properties of SiCO glasses and SiCO/Fe composites
Marco Buongiorno Nardelli	University of North Texas	1,000,000	322,899	Ab initio infrastructure for high-throughput computational materials
David N. Beratan	Duke University	400,000	0	Photoinduced Electron Transfer Between Semiconducting Nanoparticles
Robert Patton	ORNL	1,000,000	65	Modeling & Simulation of Medicare & Medicaid Services
Biswas Sengupta	Indian Institute of Science	1,323,000	202	The role of constraints in the design of the nervous system
Christian Trott	SNL-NM	200,000	124,537	LAMMPS-SNAP Titan readiness
Bhagawan Sahu	Global Foundries US Inc.	5,880,000	581,290	Density Functional Studies of Si/SiGe interface structures
John Turner	ORNL	14,700,000	7,608,936	Fundamental studies of multiphase flows and corrosion mechanisms in nuclear engineering applications
Dipangkar Dutta	Mississippi State University	1,470,000	13,686	A New Search for the Neutron Electric Dipole Moment
Kenneth Read	ORNL	347,000	48,062	Probing Fluctuating Initial Conditions of Heavy-Ion Collisions
Nitin Bangera	MIND Research Network	50,000	3	GPU Accelerated Forward Solutions for the EEG and MEG
Bobby Sumpter	ORNL	8,820,000	8,846,390	Computational Nanoscience
Misun Min	ANL	500,000	0	Nek-HOM (Codes for High Order Methods)
Hazim El-Mounayri	Indiana University	1,000,000	326,044	AFM-based nanomachining of 3D structures
Sreekanth Pannala	ORNL	1,470,000	41,211	Using Solid Particles as Heat Transfer Fluid in CSP Plants
Travis Humble	ORNL	2,000,000	0	Jade Adiabatic Device Emulator (JADE)
Cory Hauck	ORNL	2,000,000	376,335	Moment Methods for Linear Kinetic Equations
Katrin Heitmann	ANL	2,000,000	2,083,227	Dark Universe
George Vahala	College of William and Mary	2,940,000	504,127	Lattice Algorithms for Quantum and Classical Turbulence

Table D-1. 2013 Director's Discretionary Allocations (continued)

PI	Affiliation	2013 Allocation	2013 Usage	Project Name
Shanti Bhushan	Mississippi State University	1,470,000	1,005,215	Hybrid CPU/GPU Parallelization of a Pseudo-Spectral Solver for Direct Numerical Simulations of Transitional Flow
Peyman Givi	University of Pittsburg	712,812	0	US National Center for Hypersonic Combined Cycle Propulsion
Antonino Ferrante	University of Washington	1,176,000	904,253	Petascale DNS of high Reynolds number multi-phase turbulent flows
Pui-kuen Yeung	Georgia Institute of Technology	8,000,000	626,040	Scale-Similarity and Turbulence Mixing: Schmidt number effects and new algorithmic developments
Jacopo Buongiorno	Massachusetts Institute of Technology	100,000	0	CFD Simulations of Multiphase Phenomena in Pipes for Carbon Capture and Sequestration and Enhanced Oil Recovery
Andrew Corrigan	Department of Defense	100,000	0	Benchmarking the Jet Engine Noise Reduction (JENRE) code on Titan
Galen Shipman	ORNL	50,000	0	Data Intensive Science Incubators
		346,634,008	**171,404,416**	